Configuring IP Multicasting and Multimedia Services

BayRS Version 12.00
Site Manager Software Version 6.00

Part No. 117355-A Rev. A
September 1997

Bay Networks

Bay Networks

4401 Great America Parkway
Santa Clara, CA 95054

8 Federal Street
Billerica, MA 01821

Copyright © 1997 Bay Networks, Inc.

All rights reserved. Printed in the USA. September 1997.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Bay Networks, Inc.

The software described in this document is furnished under a license agreement and may only be used in accordance with the terms of that license. A summary of the Software License is included in this document.

Trademarks

ACE, AFN, AN, BCN, BLN, BN, BNX, CN, FN, FRE, GAME, LN, PPX, and Bay Networks are registered trademarks and Advanced Remote Node, ANH, ARN, ASN, Bay•SIS, BayStack, BayStream, BCNX, BLNX, IP AutoLearn, SN, SPEX, Switch Node, System 5000, Bay Networks Press, and the Bay Networks logo are trademarks of Bay Networks, Inc.

All other trademarks and registered trademarks are the property of their respective owners.

Restricted Rights Legend

Use, duplication, or disclosure by the United States Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Notwithstanding any other license agreement that may pertain to, or accompany the delivery of, this computer software, the rights of the United States Government regarding its use, reproduction, and disclosure are as set forth in the Commercial Computer Software-Restricted Rights clause at FAR 52.227-19.

Statement of Conditions

In the interest of improving internal design, operational function, and/or reliability, Bay Networks, Inc. reserves the right to make changes to the products described in this document without notice.

Bay Networks, Inc. does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Portions of the code in this software product are Copyright © 1988, Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms of such portions are permitted, provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that such portions of the software were developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from such portions of the software without specific prior written permission.

SUCH PORTIONS OF THE SOFTWARE ARE PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

In addition, the program and information contained herein are licensed only pursuant to a license agreement that contains restrictions on use and disclosure (that may incorporate by reference certain limitations and notices imposed by third parties).

Bay Networks, Inc. Software License Agreement

NOTICE: Please carefully read this license agreement before copying or using the accompanying software or installing the hardware unit with pre-enabled software (each of which is referred to as “Software” in this Agreement).

BY COPYING OR USING THE SOFTWARE, YOU ACCEPT ALL OF THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT. THE TERMS EXPRESSED IN THIS AGREEMENT ARE THE ONLY TERMS UNDER WHICH BAY NETWORKS WILL PERMIT YOU TO USE THE SOFTWARE. If you do not accept these terms and conditions, return the product, unused and in the original shipping container, within 30 days of purchase to obtain a credit for the full purchase price

1. License Grant. Bay Networks, Inc. ("Bay Networks") grants the end user of the Software ("Licensee") a personal, nonexclusive, nontransferable license: a) to use the Software either on a single computer or, if applicable, on a single authorized device identified by host ID, for which it was originally acquired; b) to copy the Software solely for backup purposes in support of authorized use of the Software; and c) to use and copy the associated user manual solely in support of authorized use of the Software by Licensee. This license applies to the Software only and does not extend to Bay Networks Agent software or other Bay Networks software products. Bay Networks Agent software or other Bay Networks software products are licensed for use under the terms of the applicable Bay Networks, Inc. Software License Agreement that accompanies such software and upon payment by the end user of the applicable license fees for such software.

2. Restrictions on use; reservation of rights. The Software and user manuals are protected under copyright laws. Bay Networks and/or its licensors retain all title and ownership in both the Software and user manuals, including any revisions made by Bay Networks or its licensors. The copyright notice must be reproduced and included with any copy of any portion of the Software or user manuals. Licensee may not modify, translate, decompile, disassemble, use for any competitive analysis, reverse engineer, distribute, or create derivative works from the Software or user manuals or any copy, in whole or in part. Except as expressly provided in this Agreement, Licensee may not copy or transfer the Software or user manuals, in whole or in part. The Software and user manuals embody Bay Networks' and its licensors' confidential and proprietary intellectual property. Licensee shall not sublicense, assign, or otherwise disclose to any third party the Software, or any information about the operation, design, performance, or implementation of the Software and user manuals that is confidential to Bay Networks and its licensors; however, Licensee may grant permission to its consultants, subcontractors, and agents to use the Software at Licensee's facility, provided they have agreed to use the Software only in accordance with the terms of this license.

3. Limited warranty. Bay Networks warrants each item of Software, as delivered by Bay Networks and properly installed and operated on Bay Networks hardware or other equipment it is originally licensed for, to function substantially as described in its accompanying user manual during its warranty period, which begins on the date Software is first shipped to Licensee. If any item of Software fails to so function during its warranty period, as the sole remedy Bay Networks will at its discretion provide a suitable fix, patch, or workaround for the problem that may be included in a future Software release. Bay Networks further warrants to Licensee that the media on which the Software is provided will be free from defects in materials and workmanship under normal use for a period of 90 days from the date Software is first shipped to Licensee. Bay Networks will replace defective media at no charge if it is returned to Bay Networks during the warranty period along with proof of the date of shipment. This warranty does not apply if the media has been damaged as a result of accident, misuse, or abuse. The Licensee assumes all responsibility for selection of the Software to achieve Licensee's intended results and for the installation, use, and results obtained from the Software. Bay Networks does not warrant a) that the functions contained in the software will meet the Licensee's requirements, b) that the Software will operate in the hardware or software combinations that the Licensee may select, c) that the operation of the Software will be uninterrupted or error free, or d) that all defects in the operation of the Software will be corrected. Bay Networks is not obligated to remedy any Software defect that cannot be reproduced with the latest Software release. These warranties do not apply to the Software if it has been (i) altered, except by Bay Networks or in accordance with its instructions; (ii) used in conjunction with another vendor's product, resulting in the defect; or (iii) damaged by improper environment, abuse, misuse, accident, or negligence. THE FOREGOING WARRANTIES AND LIMITATIONS ARE EXCLUSIVE REMEDIES AND ARE IN LIEU OF ALL OTHER WARRANTIES EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Licensee is responsible for the security of its own data and information and for maintaining adequate procedures apart from the Software to reconstruct lost or altered files, data, or programs.

4. Limitation of liability. IN NO EVENT WILL BAY NETWORKS OR ITS LICENSORS BE LIABLE FOR ANY COST OF SUBSTITUTE PROCUREMENT; SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES; OR ANY DAMAGES RESULTING FROM INACCURATE OR LOST DATA OR LOSS OF USE OR

PROFITS ARISING OUT OF OR IN CONNECTION WITH THE PERFORMANCE OF THE SOFTWARE, EVEN IF BAY NETWORKS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL THE LIABILITY OF BAY NETWORKS RELATING TO THE SOFTWARE OR THIS AGREEMENT EXCEED THE PRICE PAID TO BAY NETWORKS FOR THE SOFTWARE LICENSE.

5. Government Licensees. This provision applies to all Software and documentation acquired directly or indirectly by or on behalf of the United States Government. The Software and documentation are commercial products, licensed on the open market at market prices, and were developed entirely at private expense and without the use of any U.S. Government funds. The license to the U.S. Government is granted only with restricted rights, and use, duplication, or disclosure by the U.S. Government is subject to the restrictions set forth in subparagraph (c)(1) of the Commercial Computer Software—Restricted Rights clause of FAR 52.227-19 and the limitations set out in this license for civilian agencies, and subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause of DFARS 252.227-7013, for agencies of the Department of Defense or their successors, whichever is applicable.

6. Use of Software in the European Community. This provision applies to all Software acquired for use within the European Community. If Licensee uses the Software within a country in the European Community, the Software Directive enacted by the Council of European Communities Directive dated 14 May, 1991, will apply to the examination of the Software to facilitate interoperability. Licensee agrees to notify Bay Networks of any such intended examination of the Software and may procure support and assistance from Bay Networks.

7. Term and termination. This license is effective until terminated; however, all of the restrictions with respect to Bay Networks' copyright in the Software and user manuals will cease being effective at the date of expiration of the Bay Networks copyright; those restrictions relating to use and disclosure of Bay Networks' confidential information shall continue in effect. Licensee may terminate this license at any time. The license will automatically terminate if Licensee fails to comply with any of the terms and conditions of the license. Upon termination for any reason, Licensee will immediately destroy or return to Bay Networks the Software, user manuals, and all copies. Bay Networks is not liable to Licensee for damages in any form solely by reason of the termination of this license.

8. Export and Re-export. Licensee agrees not to export, directly or indirectly, the Software or related technical data or information without first obtaining any required export licenses or other governmental approvals. Without limiting the foregoing, Licensee, on behalf of itself and its subsidiaries and affiliates, agrees that it will not, without first obtaining all export licenses and approvals required by the U.S. Government: (i) export, re-export, transfer, or divert any such Software or technical data, or any direct product thereof, to any country to which such exports or re-exports are restricted or embargoed under United States export control laws and regulations, or to any national or resident of such restricted or embargoed countries; or (ii) provide the Software or related technical data or information to any military end user or for any military end use, including the design, development, or production of any chemical, nuclear, or biological weapons.

9. General. If any provision of this Agreement is held to be invalid or unenforceable by a court of competent jurisdiction, the remainder of the provisions of this Agreement shall remain in full force and effect. This Agreement will be governed by the laws of the state of California.

Should you have any questions concerning this Agreement, contact Bay Networks, Inc., 4401 Great America Parkway, P.O. Box 58185, Santa Clara, California 95054-8185.

LICENSEE ACKNOWLEDGES THAT LICENSEE HAS READ THIS AGREEMENT, UNDERSTANDS IT, AND AGREES TO BE BOUND BY ITS TERMS AND CONDITIONS. LICENSEE FURTHER AGREES THAT THIS AGREEMENT IS THE ENTIRE AND EXCLUSIVE AGREEMENT BETWEEN BAY NETWORKS AND LICENSEE, WHICH SUPERSEDES ALL PRIOR ORAL AND WRITTEN AGREEMENTS AND COMMUNICATIONS BETWEEN THE PARTIES PERTAINING TO THE SUBJECT MATTER OF THIS AGREEMENT. NO DIFFERENT OR ADDITIONAL TERMS WILL BE ENFORCEABLE AGAINST BAY NETWORKS UNLESS BAY NETWORKS GIVES ITS EXPRESS WRITTEN CONSENT, INCLUDING AN EXPRESS WAIVER OF THE TERMS OF THIS AGREEMENT.

Contents

About This Guide

Before You Begin xvi
Conventions xvi
Acronyms xvii
Ordering Bay Networks Publications xviii
Bay Networks Customer Service xix
How to Get Help xix

Chapter 1
Multicasting and Multimedia Overview

Multicast Host Groups 1-2
Multicast Addresses 1-3
Internet Group Management Protocol 1-3
Distance Vector Multicast Routing Protocol 1-3
Multicast Extensions to OSPF 1-4
Quality of Service Extensions to OSPF 1-4
Resource Reservation Protocol 1-4
Bay Networks Resource Manager 1-4
Multicast Table Manager 1-5

Chapter 2
Starting Multicasting and Multimedia Services

Starting IGMP 2-2
Starting DVMRP 2-3
Starting MOSPF 2-4
Starting RSVP 2-5

Chapter 3
Customizing IGMP

IGMP Concepts and Terminology 3-2
IGMP Queries 3-2
IGMP Host Reports 3-3
Host Leave Messages 3-3
Customizing Global IGMP 3-4
Enabling and Disabling IGMP 3-5
Estimating the Number of Groups 3-6
Specifying a Version Threshold Time 3-8
Configuring Logging 3-9
Enabling Join Acknowledgments 3-10
Specifying a Forwarding Cache Limit 3-11
Customizing IGMP on an Interface 3-12
Enabling and Disabling IGMP on an Interface 3-13
Specifying a Query Rate 3-14
Specifying a Membership Timeout Interval 3-16
Specifying a Designated Router Timeout Interval 3-17
Specifying a Maximum Host Response Time 3-18
Specifying the Lifetime of a Cache Entry for Mtrace 3-19
Configuring a Static Host Entry 3-20

Chapter 4
Customizing DVMRP

DVMRP Concepts and Terminology 4-2
Neighbor Connections 4-2
Source Route Advertisements 4-4
How DVMRP Chooses a Route 4-5
Routing Table 4-6
Shortest-Path Trees 4-7
Customizing DVMRP Globally 4-8
Disabling and Reenabling DVMRP 4-9
Setting the Update Interval 4-10
Configuring Tree Pruning 4-11
Configuring DVMRP Timers 4-12
Setting the Route Expiration Timer 4-12

Setting the Garbage Timer 4-13
Setting the Neighbor Report Timer 4-14
Setting the Leaf Timer 4-15
Specifying a Probe Interval 4-16
Specifying a Route Switch Interval 4-17
Estimating the Number of Routes 4-18
Logging Debugging Messages 4-19
Specifying the Maximum Number of Routes 4-20
Enabling and Disabling Routing Policies 4-21
Customizing DVMRP on a Circuit 4-22
Disabling and Reenabling DVMRP on a Circuit 4-22
Disabling Multicast Support on a Unicast Tunnel 4-23
Calculating a Route Metric and Comparing Route Costs 4-24
Specifying a Threshold 4-25
Configuring a Forwarding Table 4-27
Disabling Advertisement of Local Networks 4-29
Advertising a Default Route 4-30
Listening for a Default Route 4-31
Accepting a Report from a Nonstandard DVMRP Neighbor 4-32
Configuring a Tunnel 4-33
Supplying Addresses for the Tunnel 4-34
Disabling and Reenabling the Tunnel 4-35
Choosing the Encapsulation Mode 4-36
Specifying the Cost 4-37
Specifying a TTL 4-38
Configuring a Forwarding Table for the Tunnel 4-39
Advertising a Default Route 4-41
Listening for a Default Route 4-42
Specifying an Encapsulation Mode for Control Messages 4-43
Accepting a Report from a Nonstandard DVMRP Neighbor 4-44

Chapter 5
Customizing OSPF Multicast Extensions

MOSPF Overview 5-2
Configuring MOSPF Globally 5-4
Enabling Deterministic Multicasting 5-5

Specifying a Forwarding Timeout Value 5-6
Specifying a Maximum Number of Queued Packets 5-7
Enabling Dynamic TTL 5-8
Configuring Multicast Forwarding on an OSPF Interface 5-9

Chapter 6
Configuring QOSPF

Enabling the Deterministic Hold-Down Feature 6-3
Enabling Route Pinning 6-4
Enabling the Opaque Capability 6-5
Enabling the Opaque Capability on an Interface 6-6

Chapter 7
Configuring RSVP

How RSVP Works 7-2
Configuring RSVP Globally 7-4
Disabling and Reenabling RSVP 7-4
Choosing a Slot 7-5
Configuring RSVP Message Logging 7-6
Configuring RSVP on an Interface 7-7
Disabling and Reenabling RSVP on an Interface 7-7
Setting the RSVP Default Refresh Timer 7-8
Setting the RSVP Default Lifetime Multiplier 7-9
Setting the RSVP Refresh Blockade Multiplier 7-10
Setting the TTL Override 7-11
Setting the Route Delay Value 7-12
Enabling and Disabling UDP Encapsulation 7-13

Chapter 8
Reserving Line Resources

About Line Resource Reservation 8-2
Resource Manager 8-2
Supported Media 8-3
How the Resource Manager Works with RSVP 8-3
Setting the Estimated Bandwidth 8-4
Setting the Reservable Bandwidth 8-5
Specifying the Traffic Queuing Algorithm 8-6

Specifying the LRM Policing Algorithm 8-7
Setting the Bandwidth Interval 8-8
Setting the Inflate Reservations Percentage 8-9
Specifying the Unreserved Policing Algorithm 8-10
Specifying the Unreserved Queue Length 8-11
Specifying the Multiline Select Algorithm 8-12
Setting the Multiline Threshold Bandwidth 8-13
Setting the Reservation Latency 8-14
Setting the Maximum Reservable Bandwidth for a Dataflow 8-15
Setting the Maximum Buffer Space for a Dataflow 8-16

Appendix A
Site Manager Parameters

DVMRP Parameters A-2
DVMRP Global Parameters A-2
DVMRP Circuit Parameters A-7
DVMRP Tunnel Parameters A-11
IGMP Parameters A-15
IGMP Global Configuration Parameters A-15
IGMP Static Host Parameters A-17
IGMP Entry Interface Parameters A-18
IP Configuration Parameters A-20
Resource Manager Parameters A-22
OSPF Parameters for Multicast and QoS Extensions A-29
OSPF Global Parameters A-29
OSPF Interface Parameters A-32
RSVP Parameters A-33
RSVP Global Parameters A-33
RSVP Interface Parameters A-34

Appendix B
Multicasting Tools

Mtrace B-2
Mrinfo B-4

Index

Figures

Figure 4-1. Multicast Routers 4-3
Figure 5-1. OSPF Autonomous Systems 5-3
Figure 6-1. QOSPF Topology 6-2
Figure 7-1. RSVP Path Message 7-2
Figure 7-2. RSVP Reserve Message 7-3

Tables

Table 4-1. Parts of a Routing Table Entry 4-6
Table 4-2. Recommended TTL and Threshold Values 4-25
Table B-1. Mtrace Command Options B-2
Table B-2. Mrinfo Command Options B-4

About This Guide

If you are responsible for configuring IGMP, DVMRP, MOSPF, QOSPF, RSVP, or line services, you need to read this guide.

You can now use the Bay Command Console (BCC™) to customize IGMP on a router. In this guide, you will find instructions for using both the BCC and Site Manager for IGMP.

If you want to	Go to
Learn about IP multicasting and multimedia services	Chapter 1
Start IP multicasting and multimedia services on the router	Chapter 2
Customize IGMP	Chapter 3
Configure and customize DVMRP	Chapter 4
Configure and customize MOSPF	Chapter 5
Configure and customize QOSPF	Chapter 6
Configure and customize RSVP	Chapter 7
Reserve line resources on the router	Chapter 8
Obtain information about Site Manager parameters (this is the same information you obtain using Site Manager online Help)	Appendix A
Use multicasting tools	Appendix B

Before You Begin

Before using this guide, you must complete the following procedures. For a new router:

- Install the router (see the installation manual that came with your router).
- Connect the router to the network and create a pilot configuration file (see *Quick-Starting Routers*, *Configuring BayStack Remote Access*, or *Connecting ASN Routers to a Network)*.

Make sure that you are running the latest version of Bay Networks Site Manager and router software. For instructions, see *Upgrading Routers from Version 7–11.xx to Version 12.00*.

Conventions

angle brackets (< >)	Indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when entering the command. Example: if command syntax is **ping** *<ip_address>*, you enter **ping 192.32.10.12**
bold text	Indicates text that you need to enter, command names, and buttons in menu paths. Example: Enter **wfsm &**
	Example: Use the **dinfo** command.
	Example: ATM DXI > Interfaces > **PVCs** identifies the PVCs button in the window that appears when you select the Interfaces option from the ATM DXI menu.
brackets ([ ])	Indicate optional elements. You can choose none, one, or all of the options.
ellipsis points	Horizontal (. . .) and vertical (⋮) ellipsis points indicate omitted information.
italic text	Indicates variable values in command syntax descriptions, new terms, file and directory names, and book titles.
quotation marks (“ ”)	Indicate the title of a chapter or section within a book.

`screen text`	Indicates data that appears on the screen. Example: `Set Bay Networks Trap Monitor Filters`
separator (>)	Separates menu and option names in instructions and internal pin-to-pin wire connections. Example: Protocols > AppleTalk identifies the AppleTalk option in the Protocols menu. Example: Pin 7 > 19 > 20
vertical line (|)	Indicates that you enter only one of the parts of the command. The vertical line separates choices. Do not type the vertical line when entering the command. Example: If the command syntax is **show at routes** | **nets**, you enter either **show at routes** or **show at nets**, but not both.

Acronyms

AUI	Attachment Unit Interface
BootP	Bootstrap Protocol
BRI	Basic Rate Interface
CCITT	International Telegraph and Telephone Consultative Committee (now ITU-T)
CSMA/CD	carrier sense multiple access with collision detection
DLCMI	Data Link Control Management Interface
GUI	graphical user interface
HDLC	high-level data link control
IP	Internet Protocol
ISDN	Integrated Services Digital Network
ISO	International Organization for Standardization
ITU-T	International Telecommunications Union-Telecommunications (formerly CCITT)
LAN	local area network
MAC	media access control
MAU	media access unit
MDI-X	media-dependent interface with crossover
NBMA	nonbroadcast multi-access

OSI	Open Systems Interconnection
OSPF	Open Shortest Path First (Protocol)
PPP	Point-to-Point Protocol
SMDS	switched multimegabit data service
SNMP	Simple Network Management Protocol
STP	shielded twisted-pair
TCP/IP	Transmission Control Protocol/Internet Protocol
Telnet	Telecommunication Network
TFTP	Trivial File Transfer Protocol
TPE	twisted-pair Ethernet
UTP	unshielded twisted-pair
WAN	wide area network

Ordering Bay Networks Publications

To purchase additional copies of this document or other Bay Networks® publications, order by part number from Bay Networks Press™ at the following numbers:

- Phone--U.S./Canada: 888-422-9773
- Phone--International: 510-490-4752
- FAX--U.S./Canada and International: 510-498-2609

The Bay Networks Press catalog is available on the World Wide Web at *support.baynetworks.com/Library/GenMisc*. Bay Networks publications are available on the World Wide Web at *support.baynetworks.com/Library/tpubs*.

Bay Networks Customer Service

You can purchase a support contract from your Bay Networks distributor or authorized reseller, or directly from Bay Networks Services. For information about, or to purchase a Bay Networks service contract, either call your local Bay Networks field sales office or one of the following numbers:

Region	Telephone number	Fax number
United States and Canada	800-2LANWAN; then enter Express Routing Code (ERC) 290, when prompted, to purchase or renew a service contract 978-916-8880 (direct)	978-916-3514
Europe	33-4-92-96-69-66	33-4-92-96-69-96
Asia/Pacific	61-2-9927-8888	61-2-9927-8899
Latin America	561-988-7661	561-988-7550

Information about customer service is also available on the World Wide Web at *support.baynetworks.com*.

How to Get Help

If you purchased a service contract for your Bay Networks product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If you purchased a Bay Networks service program, call one of the following Bay Networks Technical Solutions Centers:

Technical Solutions Center	Telephone number	Fax number
Billerica, MA	800-2LANWAN	978-916-3514
Santa Clara, CA	800-2LANWAN	408-495-1188
Valbonne, France	33-4-92-96-69-68	33-4-92-96-69-98
Sydney, Australia	61-2-9927-8800	61-2-9927-8811
Tokyo, Japan	81-3-5402-0180	81-3-5402-0173

Chapter 1
Multicasting and Multimedia Overview

This overview covers the following topics:

Topic	Page
Multicast Host Groups	1-22
Multicast Addresses	1-23
Internet Group Management Protocol	1-23
Distance Vector Multicast Routing Protocol	1-23
Multicast Extensions to OSPF	1-24
Quality of Service Extensions to OSPF	1-24
Resource Reservation Protocol	1-24
Bay Networks Resource Manager	1-24
Multicast Table Manager	1-25

Multicast Host Groups

IP multicasting is a method for addressing, routing, and delivering a datagram to a collection of receivers -- called a *host group*.

Host groups can be permanent or transient.

- A *permanent host group* has a well-known, administratively assigned IP multicast group address. The address, not the membership, is permanent and defines the group. A permanent host group can consist of zero or more members.
- A *transient host group* exists only as long as it has members that need its services. IP addresses in the multicast range that are not reserved for permanent groups are available for dynamic assignment to transient host groups.

Any host system on any IP network can send a message to a multicast group using the group's IP multicast address. To receive a message addressed to a multicast group, however, the host must be a member of the group and must reside on a network where that group is registered with a local multicast router.

An IP multicasting host group can consist of zero or more members and places no restrictions on its membership. Host members can reside anywhere; they can join and leave the group at any time; and they can be members of more than one group at the same time. In order to receive a multicast message from a host group, a host must be a member of the group. However, anyone can send a multicast datagram: a host does not need to be a member of a group to send a multicast message to its members.

In general, hosts that are members of the same group reside on different networks. However, a range of multicast addresses (224.0.0.*x*) is reserved for groups that are locally scoped. All message traffic for these hosts remains on the local network. Hosts that belong to a group in this address range and that reside in different networks will not receive each other's message traffic.

Note: Multicast data packets are affected by traffic filters. Therefore, you must ensure that traffic filters configured on a multicast router do not prevent a host that is a member of a group from receiving packets intended for that group.

Multicast Addresses

Each host group is assigned a unique multicast address. To reach all members of the group, a sender uses the multicast address as the destination address of the datagram.

An IP Version 4 multicast address is a Class D address (the high-order bits are set to 1110) from 224.0.0.0 to 239.255.255.255.

The block of addresses from 224.0.0.1 to 224.0.0.255 is reserved for routing protocols and other low-level protocols. Multicast routers will not forward datagrams with addresses in this range.

Internet Group Management Protocol

The Internet Group Management Protocol (IGMP) has the following characteristics:

- Allows a host to register group memberships with the local queries router to receive any datagrams sent to this router and targeted to a group with a specific IP multicast address.
- Allows a router to learn the existence of group members on its directly attached networks. The router periodically sends a general group query message to each of its local networks. Any host that is a member of a multicasting group identifies itself by sending a response.

IGMP is described in Chapter 3.

Distance Vector Multicast Routing Protocol

DVMRP is a distance vector type of multicast routing protocol. It advertises shortest-path routes to multicasting *source networks* -- that is, any network containing hosts that have the capability to issue multicast datagrams. (In this respect, DVMRP is the opposite of RIP, which advertises routes to destination networks.)

DVMRP is described in Chapter 4.

Multicast Extensions to OSPF

Multicasting Extensions to OSPF (MOSPF) enhances the OSPF protocol by providing capabilities that allow a router to forward multicast IP traffic within an OSPF Version 2 autonomous system.

For instructions on configuring MOSPF, see Chapter 5.

Quality of Service Extensions to OSPF

QOSPF is an extension of OSPF and MOSPF that works in conjunction with the Resource Reservation Protocol to provide quality of service (QoS) for multicast data flows.

QOSPF is described in Chapter 6.

Resource Reservation Protocol

The Resource Reservation Protocol (RSVP) allows host systems in an IP network to reserve resources on RSVP-capable routers for unicast or multicast dataflows. (A dataflow is a transmission of packets from a source to one or more destinations requiring a certain QoS.)

RSVP is described in Chapter 7.

Bay Networks Resource Manager

The Bay Networks Resource Manager lets you define a certain percentage of a line's bandwidth as *reservable*. Applications that require a guaranteed quality of service can negotiate for the reservable bandwidth. When the router honors a reservation request, the Resource Manager allocates bandwidth from the reservable bandwidth, reducing the bandwidth available for other requests.

For information about the Bay Networks Resource Manager, see Chapter 8.

Multicast Table Manager

The Multicast Table Manager (MTM) makes it possible for different multicast protocols to work together on the same router. MTM does the following:

- Manages DVMRP and MOSPF
- Provides multicast routing support for RSVP
- Supports MOSPF special features
- Supports multicast tools for tracing routes
- Maintains a multicast forwarding cache
- Forwards multicast traffic

Chapter 2
Starting Multicasting and Multimedia Services

This chapter explains how to perform a basic configuration -- that is, a configuration using all available defaults -- for the IP multicast services described in this manual.

Topic	Page
Starting IGMP	2-28
Starting DVMRP	2-29
Starting MOSPF	2-30
Starting RSVP	2-31

Starting IGMP

When you configure any multicasting protocol described in this manual, Site Manager automatically configures IGMP on the slot and circuit.

IGMP is required for all types of multicasting. If you want the router to receive and forward multicast packets (that is, packets with destination addresses from 224.0.0.255 to 239.255.255.255), IGMP must be running on the slot and circuit -- even if the circuit is a point-to-point circuit that will not be involved in IGMP group queries and join messages.

Starting DVMRP

To start DVMRP on the router:

1. **Configure a circuit on a slot and connector.**
2. **Configure an IP interface on the circuit.**
3. **Add DVMRP to the IP interface.**

For information and instructions on using Site Manager to configure a circuit on a slot and connector, see *Configuring WAN Line Services* or *Configuring Ethernet, FDDI, and Token Ring Services*. When you have successfully configured the circuit, the Select Protocols window opens. Proceed as follows:

Site Manager Procedure	
You do this	**System responds**
1. In the Select Protocols window, choose the following protocols: • **IP** • **DVMRP** Then click on **OK**.	The IP Configuration window opens.
2. Set the following parameters: • **IP Address** • **Subnet Mask** • **Transmit Bcast Addr** • **UnNumbered Assoc Address** Click on **Help** or see the parameter descriptions beginning on page A-20.	
3. Click on **OK**.	Site Manager returns you to the Configuration Manager window.

When you configure DVMRP on an IP interface, DVMRP is also configured globally on the router.

DVMRP is now running on the router with default values for all global and interface parameters. You customize DVMRP by modifying DVMRP parameters. For information and instructions, see Chapter 4.

Starting MOSPF

MOSPF is a multicasting extension of OSPF. To enable MOSPF, you must add OSPF to an IP interface and enable MOSPF multicast extensions to support one of the following types of multicasting:

- Intra-area multicasting
- Intra-area and inter-area multicasting
- Intra-area and inter-AS multicasting
- All multicasting

For information about MOSPF multicasting types, see Chapter 5.

Site Manager Procedure	
You do this	**System responds**
1. Configure OSPF on the router as described in *Configuring IP Services*.	
2. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
3. Choose **IP**.	The IP menu opens.
4. Choose **OSPF**.	The OSPF menu opens.
5. Choose **Global**.	The Edit OSPF Global Parameters window opens.
6. Click on the **Multicast Extensions** parameter. Click on **Help** or see the parameter description on page A-29.	
7. Choose the type of MOSPF you want to configure, and then click on **OK**.	

Starting RSVP

To start RSVP on the router:

1. **Configure a circuit on a slot and connector.**
2. **Configure an IP interface on the circuit.**
3. **Add RSVP to the IP interface.**

For information and instructions on using Site Manager to configure a circuit on a slot and connector, see see *Configuring WAN Line Services* or *Configuring Ethernet, FDDI, and Token Ring Services*. When you have successfully configured the circuit, the Select Protocols window opens. Proceed as follows:

Site Manager Procedure	
You do this	**System responds**
1. In the Select Protocols window, choose the following protocols: • **IP** • **RSVP** Then click on **OK**.	The IP Configuration window opens.
2. Set the following parameters: • **IP Address** • **Subnet Mask** • **Transmit Bcast Addr** • **UnNumbered Assoc Address** Click on **Help** or see the parameter descriptions beginning on page A-20.	
3. Click on **OK**.	The Initial RSVP Global Parameters widow appears.
4. Click on **Save**.	Site Manager queries you to create line resources.
5. Click on **OK**.	The Edit Line Resources window opens.
6. Edit the following parameters: • **Estimated Bandwidth** • **Reservable Bandwidth** Click on **Help** or see the parameter descriptions beginning on page A-22.	Site Manager returns you to the Configuration Manager window.
7. Click on **OK**.	Site Manager returns you to the Configuration Manager window.

RSVP is now running on the router with default values for all parameters. You customize RSVP by modifying RSVP parameters. For information and instructions, see Chapter 7.

Chapter 3
Customizing IGMP

This chapter covers the following topics describing the Internet Group Management Protocol (IGMP):

Topic	Page
IGMP Concepts and Terminology	3-2
Customizing Global IGMP	3-4
Customizing IGMP on an Interface	3-12
Configuring a Static Host Entry	3-20

IGMP Concepts and Terminology

The Internet Group Management Protocol (IGMP) allows hosts to communicate their desired group memberships to their local queries router to receive any datagrams sent to this router and targeted to a group with a specific IP multicast address. Bay Networks routers support IGMP Version 2.

A router communicates with the hosts on a local network by sending IGMP queries. Hosts respond by issuing IGMP reports.

This section covers the following topics:

Topic	Page
IGMP Queries	3-2
IGMP Host Reports	3-3
Host Leave Messages	3-3

IGMP Queries

A router running IGMP periodically sends host membership queries to its attached local networks. Routers that conform to RFC 1112 generate Version 1 queries. Bay Networks routers configured with IGMP generate Version 2 queries.

If multiple IGMP routers exist on the network, Bay Networks routers designate one router to send host member questions, using the following rules:

- Choose a router that generates Version 1 queries over a router that generates Version 2 queries.
- Choose the router with the lowest IP address.

IGMP Host Reports

A host that receives a membership query from a local router can respond with a *host membership report*, one report for each joined multicast group. A host that receives a query delays its reply by a random interval and listens for a reply from any other host in the same host group. Consider a network that includes two host members -- Host A and Host B -- of the same multicast group. The router sends out a host membership query on the local network. Host A and Host B both receive the query and listen on the network for a host membership report. Host B's delay time expires first, so it responds to the query with a membership report. Hearing the response, Host A does not send a report of its own for the same group.

Each query from a router to a host includes a code field. IGMP inserts a value -- *n* -- into this field specifying the maximum time in tenths of a second within which the host must issue a reply. The host uses this value to calculate a random value between 0 and *n* tenths of a second for the period that it waits before sending a response.

If at least one host on the local network specifies that it is a member of a given group, the router will forward to that network all datagrams bearing the group's multicast address.

Upon initialization, the host may immediately issue a report for each of its supported multicast groups. The router accepts and processes these asynchronous reports the same way it accepts requested reports.

Once in a steady state, hosts and routers communicate in a way that minimizes the exchange of queries and reports.

Host Leave Messages

When a Version 2 host leaves a group, it issues a *host leave message*. The multicast router on the network issues a group-specific query to determine whether there are other group members on the network. If no host responds to the query, the router assumes that no members belonging to that group exist on that interface.

Customizing Global IGMP

When you configure a multicasting protocol on the router, IGMP is automatically enabled with default values for all global parameters. You customize global IGMP by modifying these parameters as described under the following topics:

Topic	Page
Enabling and Disabling IGMP	3-5
Estimating the Number of Groups	3-6
Specifying a Version Threshold Time	3-8
Configuring Logging	3-9
Enabling Join Acknowledgments	3-10
Specifying a Forwarding Cache Limit	3-11

Enabling and Disabling IGMP

When you configure IGMP on the router, IGMP is automatically enabled.

You can use the BCC or Site Manager to disable and reenable IGMP on the router.

Using the BCC

Navigate to the global IGMP prompt and enter

state *state*

state is

enabled
disabled

Using Site Manager

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Global**.	The IGMP Global Configuration window opens.
6. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-15.	
7. Click on **Save**.	You return to the Configuration Manager window.

Estimating the Number of Groups

By default, IGMP estimates that 20 groups will be simultaneously active on this router.

This estimate allows the router to utilize memory efficiently; exceeding this size during router operation will not cause an error but may cause the router to consume more memory than required.

Note: You do not need to include in the count any group from 224.0.0.0 to 224.0.0.255.

You can use the BCC or Site Manager to supply an estimate from 5 to 65,535 groups.

Using the BCC

Navigate to the global IGMP prompt and enter

estimated-groups *groups*

groups is an integer from 5 to 65,535.

Using Site Manager

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Global**.	The IGMP Global Configuration window opens.
6. Set the **Estimated Groups** parameter. Click on **Help** or see the parameter description on page A-16.	
7. Click on **Save**.	You return to the Configuration Manager window.

Specifying a Version Threshold Time

By default, 540 seconds can elapse after IGMP detects a Version 1 query before IGMP tries to become the designated querier.

The value you specify should be greater than the Version 1 query rate of all IGMP speakers on all networks directly connected to the router.

You can use the BCC or Site Manager to specify a version threshold time from 1 second to 65,535 seconds.

Using the BCC

Navigate to the global IGMP prompt and enter

version-threshold *threshold*

threshold is the number of seconds from 1 to 65,535.

For example, the following command sets the version threshold time to 500 seconds:

```
igmp# version-threshold 500
igmp#
```

Using Site Manager

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Global**.	The IGMP Global Configuration window opens.
6. Set the **Version Threshold Time** parameter. Click on **Help** or see the parameter description on page A-16.	
7. Click on **Save**.	You return to the Configuration Manager window.

Configuring Logging

You can use the BCC or Site Manager to specify the messages you want IGMP to log.

Using the BCC

Navigate to the global IGMP prompt and enter

debug-log-flag *flag*

flag is an integer indicating the messages you want to log.

Using Site Manager

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Global**.	The IGMP Global Configuration window opens.
6. Set the **Debug** parameter. Click on **Help** or see the parameter description on page A-16.	
7. Click on **Save**.	You return to the Configuration Manager window.

Enabling Join Acknowledgments

By default, IGMP does not respond with an acknowledgment to join requests from host groups.

Using the following Site Manager procedure, you can configure IGMP globally to respond to a join request by sending an IGMP query to the group.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Global**.	The IGMP Global Configuration window opens.
6. Set the **Join Ack Enable** parameter. Click on **Help** or see the parameter description on page A-17.	
7. Click on **Save**.	You return to the Configuration Manager window.

Note: This is a proprietary Bay Networks function. It is not part of the standard IGMP protocol.

Specifying a Forwarding Cache Limit

By default, IGMP allows a maximum of 512 entries in the MTM forwarding cache. This affects all multicasting protocols.

The MTM cache must be at least as large as individual routing protocol caches (for example, the DVMRP cache).

You can use the following Site Manager procedure to set the forwarding cache to a different value.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Global**.	The IGMP Global Configuration window opens.
6. Set the **Forward Cache Limit** parameter. Click on **Help** or see the parameter description on page A-17.	
7. Click on **Save**.	You return to the Configuration Manager window.

Customizing IGMP on an Interface

When you add a multicast protocol to an IP interface, IGMP is enabled with default values for all interface parameters. You customize IGMP on the interface by modifying these parameters as described under the following topics:

Topic	Page
Enabling and Disabling IGMP on an Interface	3-13
Specifying a Query Rate	3-14
Specifying a Membership Timeout Interval	3-16
Specifying a Designated Router Timeout Interval	3-17
Specifying a Maximum Host Response Time	3-18
Specifying the Lifetime of a Cache Entry for Mtrace	3-19

Enabling and Disabling IGMP on an Interface

When you add a multicasting protocol to an IP interface, IGMP is automatically enabled on the interface.

Note: IGMP is required for multicasting to function on the interface.

You can use the BCC or Site Manager to disable and reenable IGMP on the interface

Using the BCC

Navigate to the IGMP prompt for the interface and enter

state *state*

state is

enabled
disabled

Using Site Manager

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Entry**.	The IGMP Entry Interface Parameters window opens.
6. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-18.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a Query Rate

By default, IGMP sends a group membership query on the interface every 120 seconds.

Specifying zero disables queries only. The router still forwards multicast datagrams on this circuit.

Note: If another IGMP router on this network has taken on the query role, this router will not send out queries unless it has not heard of any queries within the number of seconds specified by the Designated Router Timeout parameter.

You can use the BCC or Site Manager to disable queries or specify a query rate from 0 to 4096 seconds.

Using the BCC

Navigate the IGMP prompt for the interface and enter

query-rate *rate*

rate is the number of seconds from 0 (disabling queries) to 4096.

For example, the following command specifies a query rate of 300 seconds:

```
igmp/2.2.2.2/3# query-rate 300
igmp/2.2.2.2/3#
```

Using Site Manager

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Entry**.	The IGMP Entry Interface Parameters window opens.
6. Set the **Interface Query Rate** parameter. Click on **Help** or see the parameter description on page A-19.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a Membership Timeout Interval

By default, a local group membership is valid for 260 seconds without the receipt of a subsequent report for that group.

You can use the BCC or Site Manager to specify an interval from 50 to 8192 seconds.

Using the BCC

Navigate to the IGMP prompt for the interface and enter

membership-timeout *seconds*

seconds is a period from 50 to 8192 seconds.

For example, the following command specifies a membership timeout period of 300 seconds:

```
igmp/2.2.2.2/3# membership-timeout 300
igmp/2.2.2.2/3#
```

Using Site Manager

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Entry**.	The IGMP Entry Interface Parameters window opens.
6. Set the **Interface Membership Timeout** parameter. Click on **Help** or see the parameter description on page A-19.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a Designated Router Timeout Interval

By default, 140 seconds can elapse after the last host query message before IGMP considers the designated router to be down.

The value you specify should be greater than the query rate of all IGMP routers on the network. The suggested value is (2 * Query Rate) + 20.

You can use the BCC or Site Manager to specify a timeout interval from 10 to 8192 seconds.

Using the BCC

Navigate to the IGMP prompt for the interface and enter

dr-timeout *seconds*

seconds is an interval from 10 to 8192 seconds.

For example, the following command sets the timeout interval to 200 seconds:

```
igmp/2.2.2.2/3# dr-timeout 200
igmp/2.2.2.2/3#
```

Using Site Manager

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Entry**.	The IGMP Entry Interface Parameters window opens.
6. Set the **Designated Router Timeout** parameter. Click on **Help** or see the parameter description on page A-19.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a Maximum Host Response Time

By default, a host must wait 100 tenths of a second before responding to a query. IGMP places this value in the code field of an IGMP query.

You can use the BCC or Site Manager to specify a value from one tenth of a second to 100 tenths of a second.

Using the BCC

Navigate to the IGMP prompt for the interface and enter

max-host-response *time*

time is an integer from 1 to 100 specifying the response time in tenths of a second.

For example, the following command sets the host response to 25 tenths of a second:

```
igmp/2.2.2.2/3# max-host-response 25
igmp/2.2.2.2/3#
```

Using Site Manager

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Entry**.	The IGMP Entry Interface Parameters window opens.
6. Set the **Max Host Response Time** parameter. Click on **Help** or see the parameter description on page A-20.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying the Lifetime of a Cache Entry for Mtrace

IGMP creates entries in the forwarding cache specifically for the Mtrace utility. By default, IGMP keeps these entries in the forwarding cache for a maximum of 30 seconds.

You can use the BCC or Site Manager to specify a different lifetime value from 30 to 8192 seconds.

Using the BCC

Navigate to the IGMP prompt for the interface and enter

mtrace-lifetime *seconds*

seconds is an integer from 30 to 8192 seconds.

For example, the following command sets the Mtrace cache entry lifetime to 60 seconds:

```
igmp/2.2.2.2/3# mtrace-lifetime 60
igmp/2.2.2.2/3#
```

Using Site Manager

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Entry**.	The IGMP Entry Interface Parameters window opens.
6. Set the **Mtrace Entry Lifetime** parameter. Click on **Help** or see the parameter description on page A-20.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring a Static Host Entry

A network may include a member of a multicast group that is unable to register its group membership with the local IGMP router, for example, a host that is a member of a multicast group but that is not running IGMP.

In this situation, you can register the group member's network statically on the router.

A registration entry includes

- A group address
- An address prefix length
- A circuit name

You can use Site Manager to supply this information.

Using Site Manager

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **IGMP**.	The IGMP menu opens.
5. Choose **Static Groups**.	The IP IGMP Static Groups window opens.
6. Set the following parameters: • **Group Address** • **Prefix Length** • **Circuit Name** Click on **Help** or see the parameter descriptions beginning on page A-17.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Chapter 4
Customizing DVMRP

This chapter describes the Distance Vector Multicast Routing Protocol (DVMRP).

Topic	Page
DVMRP Concepts and Terminology	4-2
Customizing DVMRP Globally	4-8
Customizing DVMRP on a Circuit	4-22
Configuring a Tunnel	4-33

DVMRP Concepts and Terminology

DVMRP is a multicasting protocol that provides a mechanism for routers to propagate multicast datagrams in a manner that minimizes the number of excess copies sent to any particular network.

This section covers the following topics:

Topic	Page
Neighbor Connections	4-2
Source Route Advertisements	4-4
How DVMRP Chooses a Route	4-5
Routing Table	4-6
Shortest-Path Trees	4-7

Neighbor Connections

In a DVMRP environment, *neighbors* are multicasting routers that are connected directly or through a tunnel.

- *Directly connected neighbors* are routers that have an interface to the same network.
- *Tunnel-connected neighbors* are multicast routers that communicate through a unicast network, exchanging DVRMP messages that are encapsulated in IP datagrams.

In Figure 4-1, for example, multicasting Router A has two neighbors, Router B and Router C. Router A and Router B are connected directly -- they both have interfaces to Net 6. Router A and Router C communicate through a tunnel that includes a unicast router.

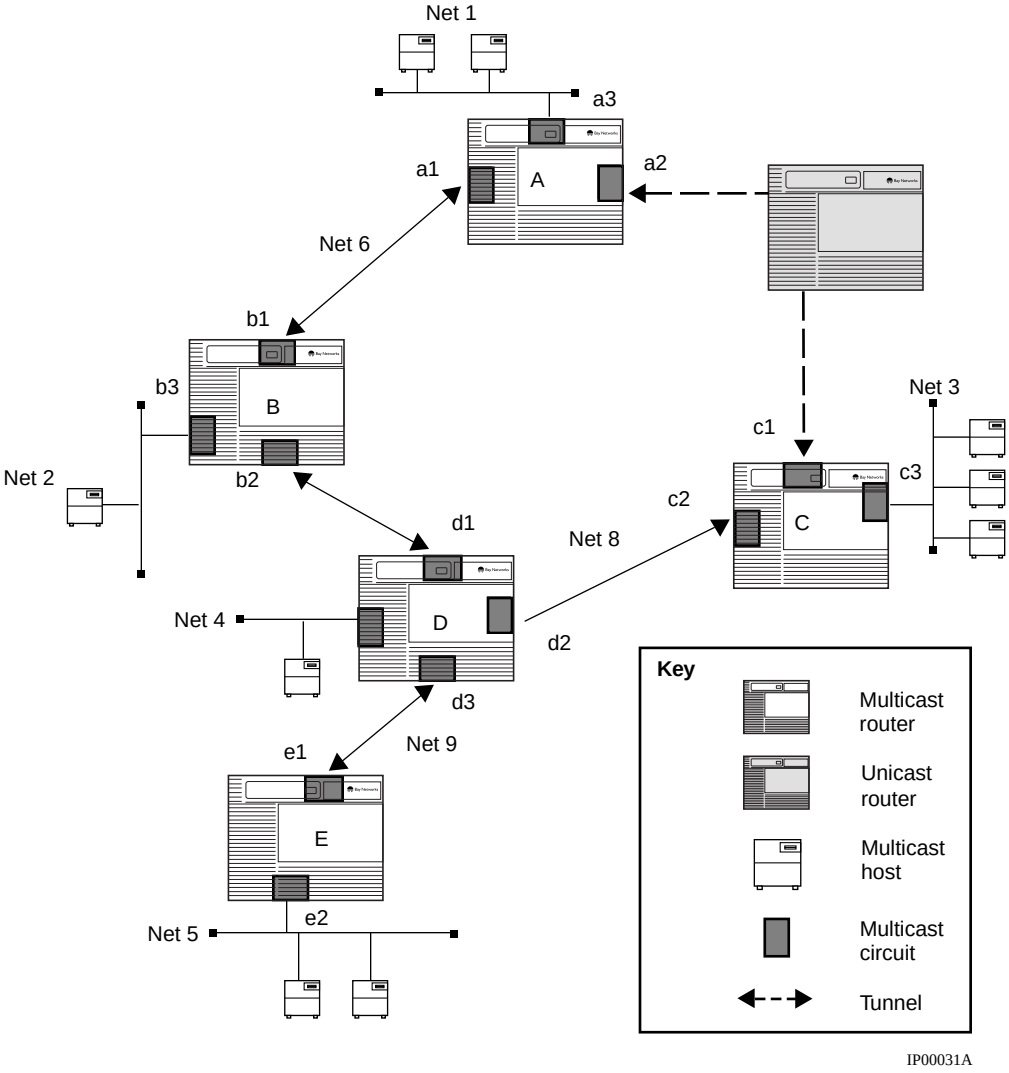

Figure 4-1. Multicast Routers

At startup, a DVMRP multicasting router

1. Initializes its routing table with information on all of its local networks
2. Learns the existence of its neighbors by sending a probe for all routes on each of its multicast interfaces (both physical circuits and tunnels)
3. Receives reports from its neighbors containing the routing information (including route costs)

In Figure 4-1, for example, Router D becomes active and issues routing probes on four multicasting interfaces. Router D receives reports from its multicasting neighbors, Routers B, C, and E.

Source Route Advertisements

A *source network* is any network containing hosts that have the capability to issue multicast datagrams. DVMRP advertises shortest-path routes to multicasting source networks. In this respect, DVMRP is the opposite of RIP, which advertises routes to destination networks.

Periodically, each multicasting router issues full or partial routing information on each DVMRP circuit, using DVMRP report messages. This routing information represents the sending router's cost to reach the specified source network. (The cost is the sum of the hop metrics along the shortest path to the given source network.)

Upon receiving a DVMRP report from another router, DVMRP reexamines its routing table to determine if the shortest path information needs updating. Specifically, DVMRP looks in the routing table for an entry describing a route to the same source network. If one exists, DVMRP compares the cost of the two routes and stores the route with the lower cost in its routing table.

A router will not send route reports on an interface until it knows (by means of received probes or reports) that it has a neighboring multicast router on that interface. It will continue to send probes periodically on an interface.

How DVMRP Chooses a Route

Each DVMRP interface is configured with a metric that indicates the cost of the hop. A router that receives multiple route reports for the same multicasting source network

- Compares the cost specified in each (based on the metric field)
- Stores information from the report with the lowest cost in its routing table

A route metric is the sum of all the interface (hop) metrics from a given route source to a given router.

In Figure 4-1 on page 4-3, for example, Router D receives two reports for the network connected to multicasting Router A, one from Router B and one from Router C. Using the metrics contained in the route reports, Router D determines that the cost of the tunneled route is greater than the cost of the route that uses direct physical connections. Router D discards the route received from Router C and stores the route received from Router B.

Router D then declares Router B to be the next-hop neighbor and Interface d1 to be the next-hop interface. After a next-hop neighbor has been declared for a route, the route updates received from that neighbor for that route take precedence until either the route times out or until another router advertises a better metric for that route.

Routing Table

Table 4-1 shows the principal items in a routing table entry.

Table 4-1. Parts of a Routing Table Entry

Item	Description
Source subnet address and mask	Address and mask of a subnetwork containing a host sourcing multicast datagrams
From-gateway	Previous-hop router leading back to the source subnet
TTL	Number of seconds before this entry is removed from the routing table

Note that the source subnet and the previous-hop router in the DVMRP routing table are the opposite of the destination subnet and next-hop router in a RIP routing table.

Using this information, the router can

1. Receive a multicast datagram and determine if the datagram has arrived on the interface that is on the shortest path to the source network.
2. Drop the datagram if it has not arrived on the shortest-path interface.
3. Flood the multicast datagram to all active downstream DVMRP neighbors.

Shortest-Path Trees

Route information used by DVMRP is independent of any other routing information used by the router -- for example, routes provided by RIP. The purpose of this routing information is to create a shortest-path tree entry in the routing table for the propagation of multicast datagrams.

The shortest-path tree entry indicates the interface that provides the shortest path to the network that is the source of the multicast datagram.

In Figure 4-1 on page 4-3, for example, the routing table on Router D includes an entry describing the shortest-path route to the network connected to Router A. The entry indicates that Interface d1 provides the shortest path to that source network.

A shortest-path tree also indicates those interfaces that are on the shortest path to that source network from a neighboring router.

Router E considers the network connecting it to Router D to be on the shortest path to the source network connected to router A. Router D has an interface (d3) that is part of the shortest path from Router E from the network connected to Router A.

If neighboring routers have the same metric to a given source network, the router with the lower IP address is responsible for propagating multicast traffic originating from that source network onto the network or tunnel that is common to these neighboring routers.

A network is considered a leaf network if it has no dependent downstream neighbors for a source. In Figure 4-1, the network connected to Router E (Net 5) is a *leaf network*.

Customizing DVMRP Globally

When you start DVMRP on the router, DVMRP runs with default values for all global protocols. You customize DVMRP by modifying parameters as described under the following topics:

Topic	Page
Disabling and Reenabling DVMRP	4-9
Setting the Update Interval	4-10
Configuring Tree Pruning	4-11
Configuring DVMRP Timers	4-12
Specifying a Probe Interval	4-16
Specifying a Route Switch Interval	4-17
Estimating the Number of Routes	4-18
Logging Debugging Messages	4-19
Specifying the Maximum Number of Routes	4-20
Enabling and Disabling Routing Policies	4-21

Disabling and Reenabling DVMRP

You can use Site Manager to disable and reenable DVMRP on the router.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-2.	
7. Click on **Save**.	You return to the Configuration Manager window.

Setting the Update Interval

Periodically, each multicasting router issues full or partial routing information on each DVMRP circuit, using DVMRP report messages. This routing information represents the sending router's cost to reach the specified network. (The cost is the sum of the hop metrics along the shortest path to the given source network.)

You can use Site Manager to

- Specify how often routing messages containing complete routing tables are sent.
- Specify the minimum amount of time between *triggered updates*. Triggered updates are sent when routing information changes. Issuing a full update restarts the triggered update timer. Therefore, the triggered update interval you specify must be shorter than the full update interval you set.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the following parameters: • **Full Update Interval** • **Triggered Update Interval** Click on **Help** or see the parameter descriptions beginning on page A-2.	
7. Click on **Save**.	You return to the Configuration Manager window.

Configuring Tree Pruning

If no members are detected on a branch of the shortest path tree for a given group, DVMRP *prunes* the branch.

The following conditions cause DVMRP to prune the shortest-path tree for a given group:

- The router determines that the attached network on which the host is located (or any other locally attached network) has no active member of the group.
- The router determines that all neighbors on all attached networks other than the next-hop network to the source have pruned off the source and group.
- The router determines that it is not located on the shortest network path to any remote network that includes a member of the group.
- The router sends a DVMRP prune message to its DVMRP neighbor. The neighbor removes the router from its forwarding table.

  If the host later rejoins the group, the router sends the neighbor a DVMRP graft message. The neighbor restores the router to its forwarding table.

Tree pruning is enabled on the router by default. You can use Site Manager to disable and reenable tree pruning on the router.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Pruning Enable** parameter. Click on **Help** or see the parameter description on page A-6.	
7. Click on **Save**.	You return to the Configuration Manager window.

Configuring DVMRP Timers

When a router adds or updates a route, it uses configurable timers to control the useful life of the route.

Note: DVMRP timers must be the same throughout the network.

Setting the Route Expiration Timer

The route expiration timer determines how long DVMRP waits for an update message indicating that a route is reachable. Upon expiration of this timer, this route is advertised as unreachable until it is refreshed or until it is garbaged.

By default, DVMRP sets the route expiration time at 200 seconds.You can use Site Manager to set the route expiration time from 20 to 4000 seconds.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Route Expiration Timeout** parameter. Click on **Help** or see the parameter description on page A-3.	
7. Click on **Save**.	You return to the Configuration Manager window.

Setting the Garbage Timer

The garbage timer determines how long DVMRP advertises a route as unreachable before it removes the route from the routing table. The difference between this value and the route expiration timeout value represents the time period that the route will be advertised as unreachable without subsequent refreshment.

By default, DVMRP sets the garbage timer to 340 seconds. You can use Site Manager to set the value from 40 to 8000 seconds.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Garbage Timeout** parameter. Click on **Help** or see the parameter description on page A-4.	
7. Click on **Save**.	You return to the Configuration Manager window.

Enter a value that is greater than the value you specified for the route expiration timer to allow enough time for the route to be advertised as unreachable.

Setting the Neighbor Report Timer

The neighbor report timer specifies how long the router waits to receive a report from a neighbor before considering the connection inactive.

By default, DVMRP waits 140 seconds. You can use Site Manager to specify a value from 40 to 8000 seconds.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Neighbor Timeout** parameter. Click on **Help** or see the parameter description on page A-3.	
7. Click on **Save**.	You return to the Configuration Manager window.

Setting the Leaf Timer

The leaf timer determines whether or not a network or tunnel local to a given interface is considered to be part of the shortest path to a given source network by any DVMRP neighbor.

If during this time the local interface has not received a route report for a given source network, this network or tunnel and its local interface are considered not to lie in the shortest path for any neighbor. In other words, the local network is not part of the shortest path to that specific source network.

By default, DVMRP sets the leaf timer at 200 seconds. You can use Site Manager to specify a value from 25 to 4000 seconds.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Leaf Timeout** parameter. Click on **Help** or see the parameter description on page A-3.	
7. Click on **Save**.	You return to the Configuration Manager window.

Specifying a Probe Interval

The probe interval determines how often DVMRP sends a probe on interfaces from which no neighbors have been heard.

By default, DVMRP sends a probe every 10 seconds. You can use Site Manager to specify a value from 5 to 30 seconds.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Neighbor Probe Interval** parameter. Click on **Help** or see the parameter description on page A-5.	
7. Click on **Save**.	You return to the Configuration Manager window.

If your neighbor is running DVMRP *mrouted*, ensure that your probe interval value matches the value used by the neighbor.

Specifying a Route Switch Interval

The route switch interval specifies how long the router should wait, without receiving a subsequent route update from the original neighbor, before switching to a different neighbor advertising equal cost for this route.

By default, DVMRP waits 140 seconds. You can use Site Manager to specify a value from 20 to 2000 seconds.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Route Switch Timeout** parameter. Click on **Help** or see the parameter description on page A-5.	
7. Click on **Save**.	You return to the Configuration Manager window.

If your neighbor is running DVMRP *mrouted*, the recommended value is 140 seconds.

Estimating the Number of Routes

By default, DVMRP preallocates a routing table for 25 entries. You can use Site Manager to specify the number of entries (10 or more).

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Estimated Routes** parameter. Click on **Help** or see the parameter description on page A-4.	
7. Click on **Save**.	You return to the Configuration Manager window.

Enter a value that the router can use for preallocating routing tables. For an MBone implementation, a value of 3000 or higher is recommended.

Note that routes are kept on a per-source-network basis, independent of multicast groups. This number must include a route for every network that is local to a circuit configured for multicasting. This allows the router to utilize memory efficiently. Exceeding this size during router operation will not cause an error but may cause the router to consume more memory than required.

Logging Debugging Messages

By default, DVMRP does not log debugging messages. You can use Site Manager to enable the logging of debugging messages and to specify the level of debug messages.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Debug Level** parameter. Click on **Help** or see the parameter description on page A-5.	
7. Click on **Save**.	You return to the Configuration Manager window.

Specifying the Maximum Number of Routes

By default, DVMRP stores all learned routes in the routing table. You can use Site Manager to specify the number of routes that DVMRP can learn per slot. This parameter limits the number of routes that can be stored in the routing table.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Max Routes** parameter. Click on **Help** or see the parameter description on page A-6.	
7. Click on **Save**.	You return to the Configuration Manager window.

If you specify a maximum number of routes, you must include a route for every local network that has a circuit configured for multicasting.

Enabling and Disabling Routing Policies

By default, routing policies are globally disabled.

This parameter allows you to configure a policy for the specified route. Bay Networks currently supports only the policies that do not advertise the specified route or route range. Bay Networks supports policies that advertise a default route for the specified route or route range.

You can use the following Site Manager procedure to enable and disable routing policies as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Global**.	The DVMRP Global Configuration window opens.
6. Set the **Policy Enable** parameter. Click on **Help** or see the parameter description on page A-6.	
7. Click on **Save**.	You return to the Configuration Manager window.

Customizing DVMRP on a Circuit

You customize DVMRP on a circuit by modifying DVMRP parameters as described under the following topics:

Topic	Page
Disabling and Reenabling DVMRP on a Circuit	4-22
Disabling Multicast Support on a Unicast Tunnel	4-23
Calculating a Route Metric and Comparing Route Costs	4-24
Specifying a Threshold	4-25
Configuring a Forwarding Table	4-27
Disabling Advertisement of Local Networks	4-29
Advertising a Default Route	4-30
Listening for a Default Route	4-31
Accepting a Report from a Nonstandard DVMRP Neighbor	4-32

Disabling and Reenabling DVMRP on a Circuit

By default, DVMRP is enabled on a circuit. You can use Site Manager to disable and reenable it.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Circuit**.	The DVMRP Circuit Parameters window opens.
6. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-7.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Disabling Multicast Support on a Unicast Tunnel

By default, DVMRP

- Propagates multicast routing information on the circuit
- Incorporates in routing updates information about the source network associated with this circuit
- Forwards multicast datagrams on this circuit in *native mode* -- that is, as multicast datagrams

Using Site Manager, you can disable multicast support on the circuit.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Circuit**.	The DVMRP Circuit Parameters window opens.
6. Set the **Route Enable** parameter. Click on **Help** or see the parameter description on page A-7.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Calculating a Route Metric and Comparing Route Costs

By default, a circuit has a cost of 1. Bay Networks recommends the following values:

Hop	Metric
LAN, or tunnel across a single LAN	1
Multihop tunnel	2 or 3
Serial link, or tunnel across a serial link	1
Backup tunnel	Primary tunnel metric + 1

You can use Site Manager to assign a cost from 1 to 31.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Circuit**.	The DVMRP Circuit Parameters window opens.
6. Set the **Metric** parameter. Click on **Help** or see the parameter description on page A-8.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a Threshold

Threshold values control the scope of datagram delivery. The threshold is the minimum IP TTL required for a multicast datagram to be forwarded out a given interface.

For *mrouted* compatibility, multicast datagrams originated by the router have a TTL of 1. These datagrams are not compared against the TTL. Unicast datagrams originated by the router (for example, for route reports issued via a tunnel) have a TTL of 255. A datagram that is to be forwarded through a tunnel is first compared against the threshold and, if accepted, is then encapsulated in an IP datagram with a TTL of 64.

Table 4-2 lists

- Originating TTL values recommended for certain types of multicast applications
- Threshold values recommended for routers to permit the forwarding of packets from these applications

These values are recommended for topologies that will connect to the MBone.

Table 4-2. Recommended TTL and Threshold Values

Multicast Application	TTL	Threshold
IETF channel 1 low-rate GSM audio	255	224
IETF channel 2 low-rate GSM audio	223	192
IETF channel 1 PCM audio	191	160
IETF channel 2 PCM audio	159	128
IETF channel 1 video	127	96
IETF channel 2 video	95	64
Local event audio	63	32
Local event video	31	1

By default, DVMRP uses a TTL value of 1. You can use Site Manager to specify a TTL value from 1 to 254.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Circuit**.	The DVMRP Circuit Parameters window opens.
6. Set the **Threshold** parameter. Click on **Help** or see the parameter description on page A-8.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring a Forwarding Table

To minimize the amount of time that it spends looking up routes, DVMRP creates and maintains a cache of frequently used routes -- called a *forwarding table* -- for each IP interface.

By default, DVMRP preallocates a forwarding table of 32 entries for each circuit. You can use Site Manager to specify the maximum number of entries (32 to 512) allowed in the forwarding table associated with this circuit.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Circuit**.	The DVMRP Circuit Parameters window opens.
6. Set the **Forward Cache Size** parameter. Click on **Help** or see the parameter description on page A-8.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

When configuring a cache size, you must also keep in mind the TTL value for entries. If you specify a large cache size, you can specify a larger TTL value. A smaller cache requires a smaller TTL for cache entries.

By default, DVMRP uses a TTL of 7200 seconds (two hours) for entries in the forwarding table. You can use Site Manager to specify a TTL from 300 to 86400 seconds.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Circuit**.	The DVMRP Circuit Parameters window opens.
6. Set the **Forward Cache TTL** parameter. Click on **Help** or see the parameter description on page A-9.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

If your neighbor is running *mrouted*, the recommended TTL value is 300 seconds.

Disabling Advertisement of Local Networks

By default, DVMRP advertises its own local networks over this circuit. Using Site Manager, you can configure DVMRP so that it does not advertise local networks over the circuit.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Circuit**.	The DVMRP Circuit Parameters window opens.
6. Set the **Advertise Self** parameter. Click on **Help** or see the parameter description on page A-9.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Advertising a Default Route

By default, DVMRP does not supply a default route on this circuit.

You can use Site Manager to enable and disable this feature as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Circuit**.	The DVMRP Circuit Parameters window opens.
6. Set the **Route Supply** parameter. Click on **Help** or see the parameter description on page A-9.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Listening for a Default Route

By default, DVMRP does not listen for a default route on this circuit.

You can use Site Manager to enable and disable this feature as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Circuit**.	The DVMRP Circuit Parameters window opens.
6. Set the **Route Listen** parameter. Click on **Help** or see the parameter description on page A-10.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Accepting a Report from a Nonstandard DVMRP Neighbor

By default, DVMRP does not accept a report from a nonstandard DVMRP neighbor.

A standard DVMRP router does not advertise a route report before sending a probe message to a neighbor. DVMRP uses the probe message to learn the neighbor's capability. Some routers, however, send only route reports; they do not send probes. If you want DVMRP to communicate with nonstandard DVMRP routers, be sure to set the Route Depend Probe parameter to Enable.

You can use Site Manager to enable and disable this feature as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Circuit**.	The DVMRP Circuit Parameters window opens.
6. Set the **Route Depend Probe** parameter. Click on **Help** or see the parameter description on page A-10.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Configuring a Tunnel

You configure a DVMRP tunnel by setting parameters as described under the following topics:

Topic	Page
Supplying Addresses for the Tunnel	4-34
Disabling and Reenabling the Tunnel	4-35
Choosing the Encapsulation Mode	4-36
Specifying the Cost	4-37
Specifying a TTL	4-38
Configuring a Forwarding Table for the Tunnel	4-39
Advertising a Default Route	4-41
Listening for a Default Route	4-42
Specifying an Encapsulation Mode for Control Messages	4-43
Accepting a Report from a Nonstandard DVMRP Neighbor	4-44

Supplying Addresses for the Tunnel

To identify a unicast tunnel, you must supply the unicast IP address of both ends of the tunnel: the local interface and the remote interface. You can use Site Manager to supply these addresses.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Tunnel**.	The DVMRP Tunnel Parameters window opens.
6. Click on **Add**.	The DVMRP Tunnel Address window opens.
7. Set the following parameters: • **Local IP Address** • **Remote IP Address** Click on **Help** or see the parameter descriptions beginning on page A-14.	
8. Click on **OK**.	You return to the DVMRP Tunnel Parameters window.

Disabling and Reenabling the Tunnel

When you configure a tunnel on a circuit, the tunnel is enabled by default. You can use Site Manager to disable and reenable the tunnel.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Tunnel**.	The DVMRP Tunnel Parameters window opens.
6. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-11.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Choosing the Encapsulation Mode

By default, DVMRP encapsulates a tunneled multicast datagram in an IP unicast datagram. Using Site Manager, you can configure DVMRP to loosely encapsulate multicast datagrams using the LSSR option.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Tunnel**.	The DVMRP Tunnel Parameters window opens.
6. Set the **Encapsulation Mode** parameter. Click on **Help** or see the parameter description on page A-11.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

See RFC 1075 for information about the LSSR option, which is provided for backward compatibility.

Specifying the Cost

By default, DVMRP assigns a metric of 1 to a tunnel. You can use Site Manager to specify a cost value from 1 to 31.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Tunnel**.	The DVMRP Tunnel Parameters window opens.
6. Set the **Metric** parameter. Click on **Help** or see the parameter description on page A-11.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying a TTL

By default, DVMRP assigns a TTL value of 1 to a tunnel. This value is the minimum IP TTL required for a multicast datagram to be forwarded out this tunnel. You can use Site Manager to assign a value from 1 to 254.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Tunnel**.	The DVMRP Tunnel Parameters window opens.
6. Set the **Threshold** parameter. Click on **Help** or see the parameter description on page A-12.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Use this value to control the scope of the datagrams. If the IP TTL is less than the threshold value you specify, the router drops the datagram.

Configuring a Forwarding Table for the Tunnel

By default, DVMRP preallocates a forwarding table for 32 entries. You can use Site Manager to specify the maximum number of entries (from 32 to 512) allowed in the forwarding table associated with this tunnel interface.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Tunnel**.	The DVMRP Tunnel Parameters window opens.
6. Set the **Forward Cache Size** parameter. Click on **Help** or see the parameter description on page A-12.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

When configuring a cache size, you must also keep in mind the TTL value for entries. If you specify a large cache size, you can specify a larger TTL value. A smaller cache requires a smaller TTL for cache entries.

By default, DVMRP assigns a TTL value of 7200 seconds to an entry in a tunnel forwarding table. You can use Site Manager to assign a value from 300 to 86400 seconds.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Tunnel**.	The DVMRP Tunnel Parameters window opens.
6. Set the **Forward Cache TTL** parameter. Click on **Help** or see the parameter description on page A-12.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Advertising a Default Route

By default, DVMRP does not supply a default route on this tunnel.

You can use Site Manager to enable and disable this feature as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Tunnel**.	The DVMRP Tunnel Parameters window opens.
6. Set the **Route Supply** parameter. Click on **Help** or see the parameter description on page A-13.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Listening for a Default Route

By default, DVMRP does not listen for a default route on this tunnel.

You can use Site Manager to enable and disable this feature as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Tunnel**.	The DVMRP Tunnel Parameters window opens.
6. Set the **Route Listen** parameter. Click on **Help** or see the parameter description on page A-13.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Specifying an Encapsulation Mode for Control Messages

To maintain backward compatibility with earlier versions of DVMRP, DVMRP allows you configure the encapsulation mode for DVMRP control packets.

You can use the following Site Manager procedure to perform this operation.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Tunnel**.	The DVMRP Tunnel Parameters window opens.
6. Set the **Control Message Mode** parameter. Click on **Help** or see the parameter description on page A-13.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Accepting a Report from a Nonstandard DVMRP Neighbor

By default, DVMRP does not accept a report from a nonstandard DVMRP neighbor.

A standard DVMRP router does not advertise a route report before it has sent a probe message to a neighbor. DVMRP uses the probe message to learn the neighbor's capability. Some routers, however, send only route reports; they do not send probes. If you want DVMRP to communicate with nonstandard DVMRP routers, be sure to set the Route Depend Probe parameter to Enable.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **Multicast**.	The Multicast menu opens.
4. Choose **DVMRP**.	The DVMRP menu opens.
5. Choose **Tunnel**.	The DVMRP Tunnel Parameters window opens.
6. Set the **Report Depend Probe** parameter. Click on **Help** or see the parameter description on page A-14.	
7. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Chapter 5
Customizing OSPF Multicast Extensions

MOSPF is a multicasting extension of OSPF. You configure MOSPF by setting OSPF global parameters, interface parameters, and area parameters. This chapter covers the following topics:

Topic	Page
MOSPF Overview	5-2
Configuring MOSPF Globally	5-4
Configuring Multicast Forwarding on an OSPF Interface	5-9

MOSPF Overview

MOSPF is a multicasting extension of the OSPF protocol that allows a router to forward multicast IP traffic within an OSPF Version 2 AS.

An OSPF environment typically consists of an AS divided into multiple areas connected by an OSPF backbone. You can configure an OSPF/MOSPF router as a

- Multicasting *internal router* (IR) that establishes neighbor relationships with adjacent routers within an area
- Multicasting *border router* (BR) with interfaces to one or more areas and to the OSPF backbone
- Multicasting *boundary router* (ASBR) with one or more interfaces to external autonomous systems

Figure 5-1 shows an OSPF/MOSPF AS consisting of two OSPF areas connected by an OSPF backbone.

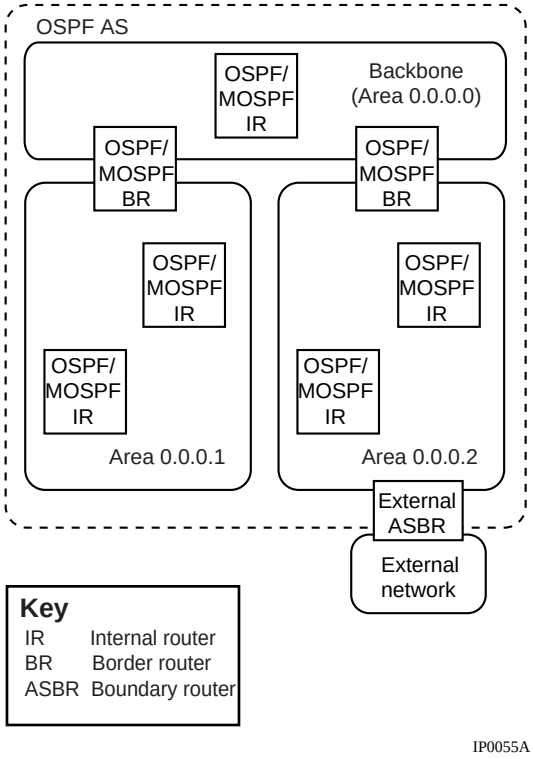

Figure 5-1. OSPF Autonomous Systems

MOSPF supports four types of multicasting:

- Intra-area multicasting for OSPF internal routers
- Intra-area and inter-area multicasting for OSPF border routers
- Intra-area and inter-AS multicasting for OSPF boundary routers
- All multicasting types

Configuring MOSPF Globally

You can use Site Manager to add MOSPF extensions to OSPF.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF**.	The OSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **Multicast Extensions** parameter. Click on **Help** or see the parameter description on page A-29.	
6. Click on **OK**.	You return to the Configuration Manager window.

After you configure OSPF with multicast extensions, you can set multicast global parameters as described under the following topics:

Topic	Page
Enabling Deterministic Multicasting	5-5
Specifying a Forwarding Timeout Value	5-6
Specifying a Maximum Number of Queued Packets	5-7
Enabling Dynamic TTL	5-8

Note: MOSPF will not forward datagrams out an interface until multicast forwarding is enabled on the interface. For instructions, refer to “Configuring Multicast Forwarding on an OSPF Interface” on page 5-9.

Enabling Deterministic Multicasting

MOSPF runs in two modes: nondeterministic mode and deterministic mode.

- In *nondeterministic mode*, MOSPF selects paths with the lowest metric cost.
- In *deterministic mode*, MOSPF uses OSPF quality of service extensions (called QOSPF) to choose paths that provide the requested QoS.

By default, deterministic MOSPF is disabled. You can use Site Manager to enable deterministic MOSPF globally.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF**.	The OSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **Multicast Deterministic** parameter. Click on **Help** or see the parameter description on page A-29.	
6. Click on **OK**.	You return to the Configuration Manager window.

When you choose MOSPF deterministic mode, you are enabling OSPF with default values for all QOSPF parameters. You customize QOSPF by modifying OSPF parameters as described in Chapter 6.

Specifying a Forwarding Timeout Value

You can use Site Manager to specify a value for timing out MOSPF entries in the forwarding table.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF**.	The OSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **Timeout Value** parameter. Click on **Help** or see the parameter description on page A-30.	
6. Click on **OK**.	You return to the Configuration Manager window.

Specifying a Maximum Number of Queued Packets

By default, MOSPF queues a maximum of 64 packets per slot for all dataflows.

You can control how many multicast packets can be queued on each slot for all flows before MOSPF finishes building multicast trees.

If you disable queuing, packets are dropped before the tree is calculated.

If you specify a maximum value, the first few data packets for a multicast flow are queued before MOSPF finishes building the multicast tree.

You can use the following Site Manager procedure to disable queuing or specify a different maximum number of queued packets.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF**.	The OSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **Max Queued Multcast Pkts** parameter. Click on **Help** or see the parameter description on page A-31.	
6. Click on **OK**.	You return to the Configuration Manager window.

Enabling Dynamic TTL

You can configure MOSPF to drop multicast packets that do not have a TTL high enough to reach the closest group member.

By default, this feature is disabled.

You can use the following Site Manager procedure to enable and disable the feature as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF**.	The OSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **Dynamic TTL** parameter. Click on **Help** or see the parameter description on page A-31.	
6. Click on **OK**.	You return to the Configuration Manager window.

Configuring Multicast Forwarding on an OSPF Interface

By default, multicast packets are forwarded on an OSPF interface. You can use Site Manager to configure OSPF forwarding on the interface as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF**.	The OSPF menu opens.
4. Choose **Interfaces**.	The Edit OSPF Interfaces window opens.
5. Set the **Multicast Forwarding** parameter. Click on **Help** or see the parameter description on page A-32.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Chapter 6
Configuring QOSPF

QOSPF is an extension of OSPF that works in conjunction with RSVP to provide quality of service for multicast data flows.

For OSPF, the best path is the path with the least metric cost. OSPF does not consider network resource information when it selects the best path to a destination.

With QOSPF extensions added, OSPF/MOSPF calculates QoS routes that can provide the best resources for a multicast flow even though the route may not be the path with the least metric cost.

In Figure 6-1, for example, Host 3 (H3) on Network 2 wants to send multicast data at rate R to a group whose members (H1 and H2) reside on Network 1. Without QOSPF extensions, MOSPF in Router 3 forwards the data to Router 1 (using the interface with lowest metric cost), whether or not Router 1 has the resources to transmit data at rate R.

Now assume that Router 2 has the resources to transmit data at rate R and Router 1 does not. With QOSPF enabled on all routers, OSPF forwards the data to Router 2, even though the interface has a higher metric cost.

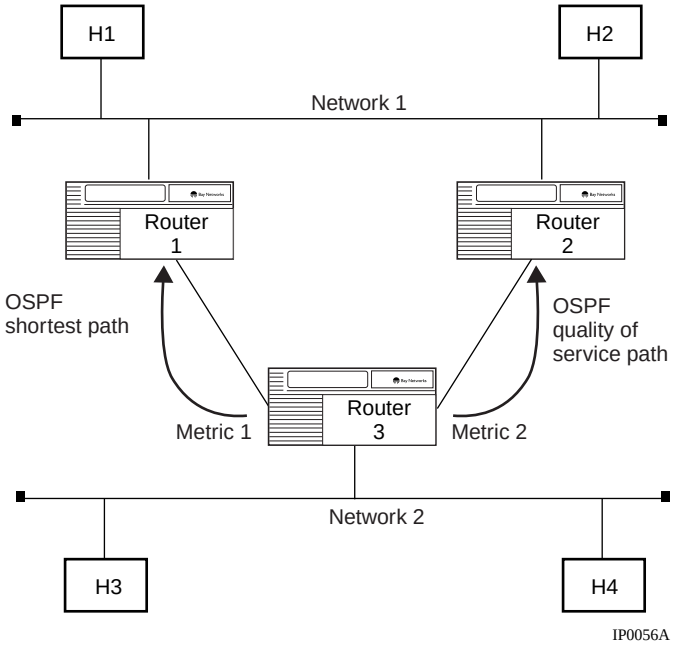

Figure 6-1. QOSPF Topology

You add QOSPF extensions to OSPF by enabling MOSPF in deterministic mode as described in “Enabling Deterministic Multicasting” on page 5-5.

When you select deterministic mode, QOSPF runs with default values for global and interface parameters. You customize QOSPF by modifying these parameters as described under the following topics:

Topic	Page
Enabling the Deterministic Hold-Down Feature	6-3
Enabling Route Pinning	6-4
Enabling the Opaque Capability	6-5
Enabling the Opaque Capability on an Interface	6-6

Enabling the Deterministic Hold-Down Feature

The deterministic hold-down feature ensures that a data flow does not go out an interface unless there is a reservation for the flow on the interface.

By default the hold-down feature is disabled. You can use Site Manager to enable this feature.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF**.	The OSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **Deterministic Mcast Hold Down** parameter. Click on **Help** or see the parameter description on page A-30.	
6. Click on **OK**.	You return to the Configuration Manager window.

Enabling Route Pinning

Route pinning ensures that once resource reservations have been established on a route from a source network to a particular receiver, the router will not replace that route with a better route (unless the established route is no longer usable).

By default, route pinning is disabled. If you have enabled MOSPF globally, you can use Site Manager to enable route pinning.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF**.	The OSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **Multicast Route Pinning** parameter. Click on **Help** or see the parameter description on page A-29.	
6. Click on **OK**.	You return to the Configuration Manager window.

Enabling the Opaque Capability

The OSPF opaque capability provides support for additional types of LSAs and additional extensions. QOSPF requires the opaque capability.

By default, the opaque capability is disabled.

You can use Site Manager to enable and disable this feature globally as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF**.	The OSPF menu opens.
4. Choose **Global**.	The Edit OSPF Global Parameters window opens.
5. Set the **Opaque Capability** parameter. Click on **Help** or see the parameter description on page A-30.	
6. Click on **OK**.	You return to the Configuration Manager window.

Enabling the Opaque Capability on an Interface

By default, if you have enabled the OSPF opaque capability globally, the feature is turned on for each OSPF interface. You can use Site Manager to set the state of the opaque capability as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **OSPF**.	The OSPF menu opens.
4. Choose **Interfaces**.	The Edit OSPF Interfaces window opens.
5. Set the **Opaque On** parameter. Click on **Help** or see the parameter description on page A-32.	
6. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Chapter 7
Configuring RSVP

The Resource Reservation Protocol (RSVP) allows host systems in an IP network to reserve resources on RSVP-capable routers for

- Unicast dataflows
- Multicast dataflows.

(A dataflow is a transmission of packets from a source to one or more destinations requiring a specific QoS.)

This chapter covers the following topics:

Topic	Page
How RSVP Works	7-2
Configuring RSVP Globally	7-4
Configuring RSVP on an Interface	7-7

How RSVP Works

A network host that wants to establish a dataflow with a specified QoS sends an RSVP path message addressed to a potential receiving host or group of hosts. In Figure 7-1, for example, a host sends a multicast RSVP path message to a group of four hosts. Routers in the network receive the multicast RSVP messages and forward them to the next-hop router until the messages reach their host destinations. The sequence of networks and routers that a path message traverses from source to destination establishes the path for the proposed dataflow.

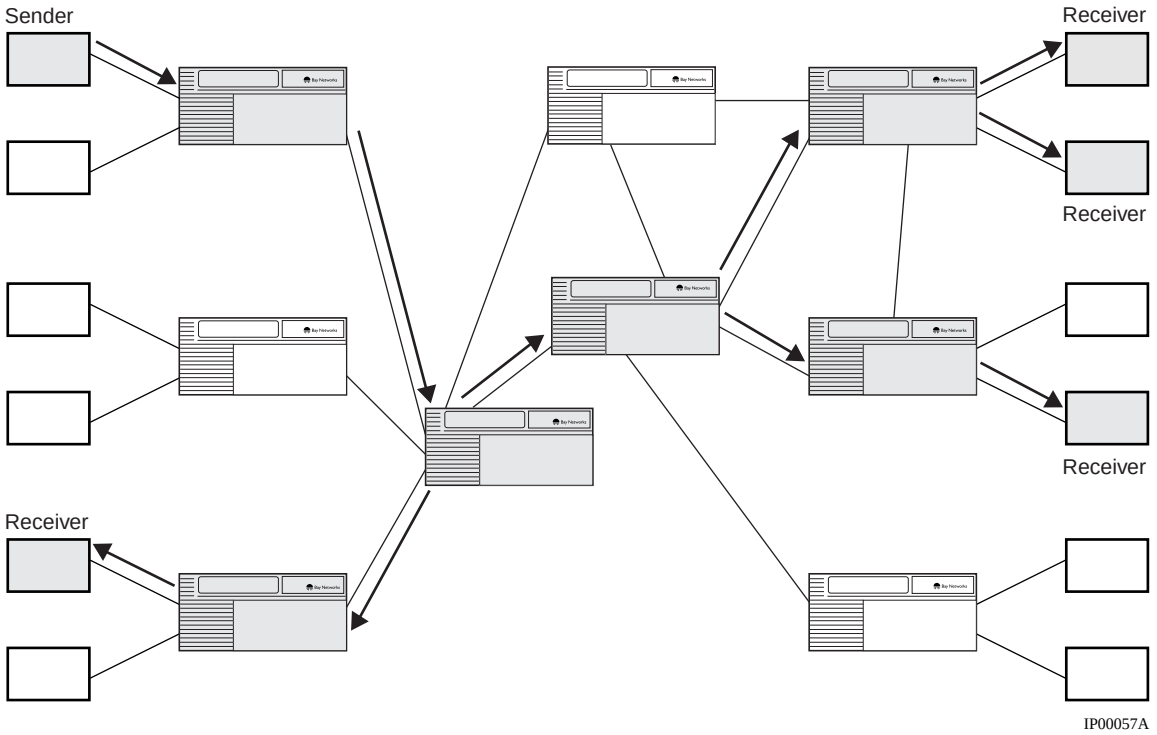

Figure 7-1. RSVP Path Message

A network host that receives an RSVP path message decides whether it wants to receive data on the proposed dataflow. If so, the host returns an RSVP reserve message. In Figure 7-2, for example, each host that has received an RSVP path message returns a reserve message. Each router in the network

1. Passes the reservation request to the local Circuit Resource Manager (CRM).
2. Forwards the reservation request to the previous hop in the flow (as determined by the earlier path message).

In this manner, the reserve message traces the data path set up by the corresponding path message and establishes resource reservation along the data path at each router. Note that reservations from multiple receivers may be merged on their way to the sender of the path message.

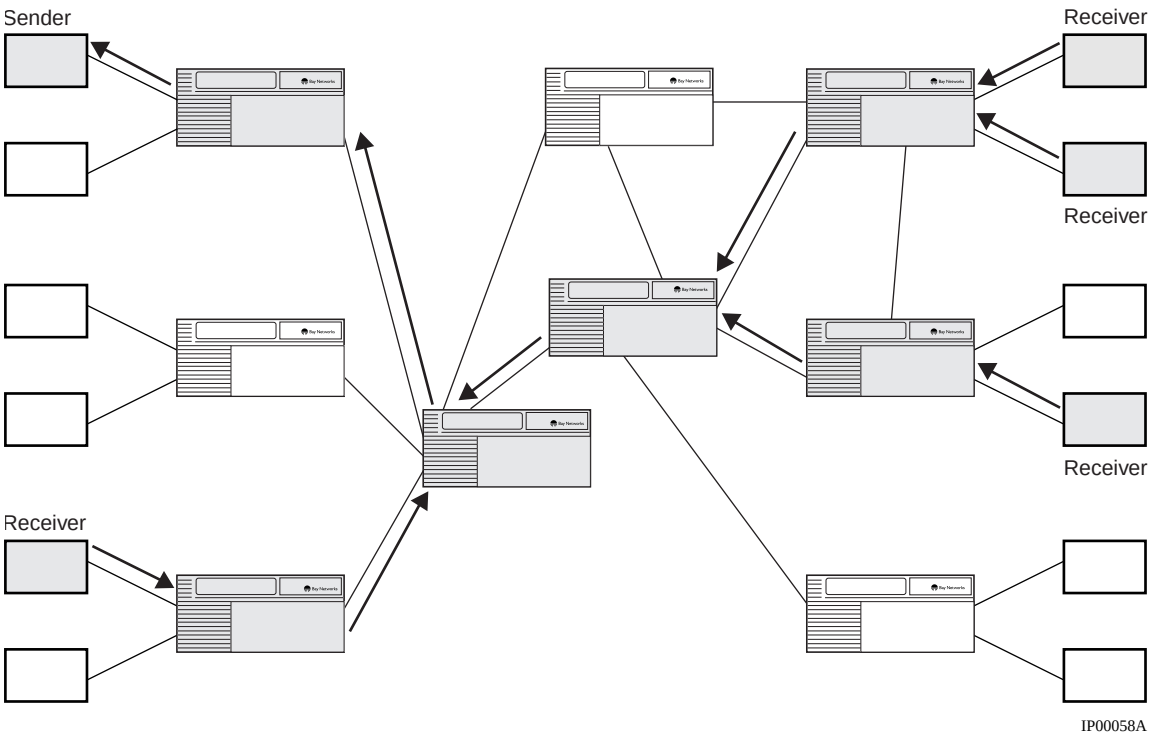

Figure 7-2. RSVP Reserve Message

Configuring RSVP Globally

When you start RSVP on the router, RSVP runs with default values for all global protocols. You customize RSVP by modifying parameters as described under the following topics:

Topic	Page
Disabling and Reenabling RSVP	7-4
Choosing a Slot	7-5
Configuring RSVP Message Logging	7-6

Disabling and Reenabling RSVP

By default, RSVP is enabled on the router.

You can use Site Manager to set the global state of RSVP as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RSVP Global**.	The RSVP Base Group Record window opens.
4. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-33.	
5. Click on **Save**.	You return to the Configuration Manager window.

Choosing a Slot

By default, RSVP runs on every slot where it is configured.

You can use Site Manger to specify the slot or slots on which you want RSVP to run.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RSVP Global**.	The RSVP Base Group Record window opens.
4. Set the **Soloist Slots** parameter. Click on **Help** or see the parameter description on page A-33.	
5. Click on **Save**.	You return to the Configuration Manager window.

Configuring RSVP Message Logging

You can use Site Manager to specify the messages you want to log.

Caution: Enabling message logging increases the chance of losing information because of lack of memory.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RSVP Global**.	The RSVP Base Group Record window opens.
4. Set the **Log Filter** parameter. Click on **Help** or see the parameter description on page A-33.	
5. Click on **Save**.	You return to the Configuration Manager window.

Configuring RSVP on an Interface

When you add RSVP to an IP interface, RSVP is configured on the interface with default values for all interface parameters. You customize RSVP on the interface by setting interface parameters as described under the following topics:

Topic	Page
Disabling and Reenabling RSVP on an Interface	7-7
Setting the RSVP Default Refresh Timer	7-8
Setting the RSVP Default Lifetime Multiplier	7-9
Setting the RSVP Refresh Blockade Multiplier	7-10
Setting the TTL Override	7-11
Setting the Route Delay Value	7-12
Enabling and Disabling UDP Encapsulation	7-13

Disabling and Reenabling RSVP on an Interface

By default, RSVP is enabled on every interface on which it is configured. You can use Site Manager to change the state of RSVP on an interface as required.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RSVP Circuit**.	The Edit RSVP Circuit Parameters window opens.
4. Set the **Enable** parameter. Click on **Help** or see the parameter description on page A-34.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Setting the RSVP Default Refresh Timer

RSVP sends periodic refresh messages to maintain the state along a reserved path. Without periodic refreshment, the state automatically times out.

The refresh timer interval is expressed in hundredths of a second. By default, the RSVP Refresh Timer is set to 3000 hundredths of a second (30 seconds).

You can use Site Manager to set the Refresh Timer for an interface. You must express the interval in hundredths of a second.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RSVP Circuit.**	The Edit RSVP Circuit Parameters window opens.
4. Set the **Refresh Interval** parameter. Click on **Help** or see the parameter description on page A-35.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Setting the RSVP Default Lifetime Multiplier

The Lifetime Multiplier is an integer that multiplies the TIME_VALUE object specified in a path or reserve message. The resulting value indicates the amount of time that the router will hold the path or reserve state without receiving a refresh message for that state.

By default, the Lifetime Multiplier is 3. You can use Site Manager to specify a multiplier.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RSVP Circuit**.	The Edit RSVP Circuit Parameters window opens.
4. Set the **Refresh Multiple** parameter. Click on **Help** or see the parameter description on page A-34.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Setting the RSVP Refresh Blockade Multiplier

By default, an RSVP interface in the blockade state will stay alive for three refresh intervals.

You can use Site Manager to set the refresh blockade multiplier.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RSVP Circuit**.	The Edit RSVP Circuit Parameters window opens.
4. Set the **Refresh Blockade Multiple** parameter. Click on **Help** or see the parameter description on page A-34.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Setting the TTL Override

By default, RSVP uses the IP TTL value. You can use Site Manager to override this value.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RSVP Circuit**.	The Edit RSVP Circuit Parameters window opens.
4. Set the **TTL** parameter. Click on **Help** or see the parameter description on page A-35.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Setting the Route Delay Value

The route delay value specifies the approximate period that elapses between the time a route is changed to the time a resulting message appears on the interface.

The route delay is expressed in hundredths of a second. By default, the route delay is 200 hundredths of a second (2 seconds). You can use Site Manager to specify a different delay. You must express the delay in hundredths of a second.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RSVP Circuit**.	The Edit RSVP Circuit Parameters window opens.
4. Set the **Route Delay** parameter. Click on **Help** or see the parameter description on page A-35.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Enabling and Disabling UDP Encapsulation

By default, RSVP uses UDP encapsulation on this interface only if it determines that a neighbor on the interface uses UDP encapsulation. Enable this feature on the interface if manual configuration requires UDP encapsulation.

You can use Site Manager to enable and disable UDP encapsulation on the interface.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, choose **Protocols**.	The Protocols menu opens.
2. Choose **IP**.	The IP menu opens.
3. Choose **RSVP Circuit**.	The Edit RSVP Circuit Parameters window opens.
4. Set the **Udp Required** parameter. Click on **Help** or see the parameter description on page A-36.	
5. Click on **Apply**, and then click on **Done**.	You return to the Configuration Manager window.

Chapter 8
Reserving Line Resources

This chapter covers the following topics:

Topic	Page
About Line Resource Reservation	8-2
Setting the Estimated Bandwidth	8-4
Setting the Reservable Bandwidth	8-5
Specifying the Traffic Queuing Algorithm	8-6
Specifying the LRM Policing Algorithm	8-7
Setting the Bandwidth Interval	8-8
Setting the Inflate Reservations Percentage	8-9
Specifying the Unreserved Policing Algorithm	8-10
Specifying the Unreserved Queue Length	8-11
Specifying the Multiline Select Algorithm	8-12
Setting the Multiline Threshold Bandwidth	8-13
Setting the Reservation Latency	8-14
Setting the Maximum Reservable Bandwidth for a Dataflow	8-15
Setting the Maximum Buffer Space for a Dataflow	8-16

About Line Resource Reservation

Emerging real-time multicast applications for digitized voice and video (such as multimedia conferencing and virtual reality) require a level of consistent network service that routers cannot meet using simple first-in-first-out (FIFO) queuing mechanisms. The existing best-effort delivery service of IP networks, where variable queuing delays and data loss due to congestion are acceptable, is not sufficient.

Resource reservation protocols such as RSVP address the requirements of delay-sensitive applications by distributing information among routers to achieve a guaranteed quality of service (QoS) for specific packet streams (flows).

By reserving bandwidth for QoS requests, a system of line resource reservation can meet the demands of real-time applications and maintain bandwidth for other traffic.

Bay Networks line resource reservation supports bandwidth guarantees only. It does not support delay guarantees.

Resource Manager

The Bay Networks Resource Manager lets you define a certain percentage of a line's bandwidth as *reservable*. Applications that require a guaranteed QoS can negotiate for the reservable bandwidth. When the router honors a reservation request, the Resource Manager allocates bandwidth from the reservable bandwidth, reducing the bandwidth available for other requests.

Supported Media

You must enable the Resource Manager on 10-Mb/s and 100-Mb/s Ethernet, FDDI, MCT1, HSSI, and synchronous interfaces that are configured with RSVP. The Resource Manager supports the following synchronous interface encapsulation methods:

- Point-to-Point Protocol (PPP)
- Bay Networks Standard
- FDDI
- 100 Mb/s Ethernet

Note: You cannot use line resource management on a line you configured for PPP multilink.

How the Resource Manager Works with RSVP

The exchange between RSVP and the Resource Manager consists of the following steps:

1. The router receives an RSVP reserve message. The reserve message includes a flow specification (*flowspec*) that indicates the resources needed to meet the QoS.
2. RSVP passes the flowspec to the Resource Manager.
3. The Resource Manager compares the requested resources with available reservable bandwidth. The Resource Manager determines reservable bandwidth from configured values described later in this chapter. It also accounts for any active reserved flows.
4. The Resource Manager either reserves the necessary services or refuses service.

Setting the Estimated Bandwidth

To activate resource reservation, specify a line and enter the portion of the line's bandwidth to make available for RSVP requests.

Enter an estimated total bandwidth for this line, from 0 to 214748364 b/s. To enable line resource management, enter a value greater than 0.

For point-to-point lines, you can usually enter the total line speed.

For an Ethernet line, you must estimate a value, because shared-media lines may not have all of the bandwidth available. Base the estimate on line statistics, the number of connected workstations, or other site-specific information.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Estimated Bandwidth** parameter. Click on **Help** or see the parameter description on page A-22.	
4. Click on **OK**.	You return to the Configuration Manager window.

Setting the Reservable Bandwidth

To enable RSVP reservable resources, specify the portion of this line's bandwidth to make available for RSVP requests for guaranteed service. The value can be any integer from 0 to 214748364 b/s, and must be less than the value of the Estimated Bandwidth parameter.

Caution: Never make all available bandwidth reservable. Leave at least 15 percent of bandwidth available for network management, routing protocols, and other best-effort traffic.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Reservable Bandwidth** parameter. Click on **Help** or see the parameter description on page A-23.	
4. Click on **OK**.	You return to the Configuration Manager window.

Specifying the Traffic Queuing Algorithm

By default, the Line Resource Manager (LRM) uses best-effort scheduling for all reserved traffic.

You can select a priority queuing algorithm to schedule reserved traffic ahead of unreserved traffic.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Traffic Queuing Algorithm** parameter. Click on **Help** or see the parameter description on page A-23.	
4. Click on **OK**.	You return to the Configuration Manager window.

Specifying the LRM Policing Algorithm

By default, LRM does not check RSVP packets against the flowspec (policing). This assumes that an upstream router is policing traffic, or the applications generating the reserved traffic consistently adhere to the flowspec and do not require policing.

When using the leaky bucket policing algorithm for reserved traffic, LRM makes sure that all packets using reserved bandwidth follow the flowspec that appears in the reserve message. For the controlled load service, LRM attempts to carry any packets that do not adhere to the flowspec.

Note that leaky bucket policing requires additional processing by the router.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Policing Algorithm** parameter. Click on **Help** or see the parameter description on page A-24.	
4. Click on **OK**.	You return to the Configuration Manager window.

Setting the Bandwidth Interval

By default, LRM measures instantaneous bandwidth over a 10-second interval. If 10 seconds is not sufficient, you can specify a different bandwidth interval for this interface, from 1 to 214748364 seconds.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Bandwidth Interval (Secs**) parameter. Click on **Help** or see the parameter description on page A-24.	
4. Click on **OK**.	You return to the Configuration Manager window.

Setting the Inflate Reservations Percentage

By default, LRM does not add a safety tolerance to each RSVP reservation on a line. To inflate (or overallocate) reservations made on this line, enter a percent value to increase each bandwidth request.

Increase the value of this parameter if you notice that LRM is discarding packets because applications generating reserved traffic are exceeding their flowspecs.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Inflate Reservations Percentage** parameter. Click on **Help** or see the parameter description on page A-25.	
4. Click on **OK**.	You return to the Configuration Manager window.

Specifying the Unreserved Policing Algorithm

LRM uses one of two policing algorithms for unreserved traffic:

- *Queue limit* -- Restricts the number of buffers (packets) of unreserved traffic that LRM queues

  This method of policing enables unreserved traffic to use available reserved bandwidth.

Note: If you enable priority queuing for this line, the router uses the priority queue parameters to perform queue limit policing. See *Configuring Traffic Filters and Protocol Prioritization* for information about priority queue parameters.

- *Leaky bucket* -- Causes the router to actively police the unreserved traffic based on the unreserved bandwidth

  This method does not allow unreserved traffic to take advantage of available reserved bandwidth.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Un reserved Policing Algorithm** parameter. Click on **Help** or see the parameter description on page A-25.	
4. Click on **OK**.	You return to the Configuration Manager window.

Specifying the Unreserved Queue Length

If the Unreserved Policing Algorithm parameter is set to Queue Limit, you can specify the maximum number of unreserved (best-effort) packets to be held in queue for transmission. After the queue length reaches this value, the router discards best-effort traffic when congestion occurs.

Note that priority queuing limits, if configured, override the value of this parameter.

By default, LRM holds a maximum of 20 unreserved packets in queue for transmission. Change this value to adjust the queue length limit from 0 to 214748364 packets (limited by available buffers).

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Unreserved Queue Length** parameter. Click on **Help** or see the parameter description on page A-26.	
4. Click on **OK**.	You return to the Configuration Manager window.

Specifying the Multiline Select Algorithm

If you configured LRM on a multiline circuit, you can choose one of two algorithms that determine how LRM selects which line to use for a new resource request:

- *First fit* -- Always uses the first available line that can service reserved traffic requests
- *Round robin* -- Rotates the use of lines available to service requests

All lines for a circuit must use the same algorithm. If any one line on a circuit specifies First Fit, all lines use the first fit algorithm.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Multiline Select Algorithm**[a] parameter**.** Click on **Help** or see the parameter description on page A-26.	
4. Click on **OK**.	You return to the Configuration Manager window.

a. You set this parameter only if you configured the Resource Manager on a multiline circuit.

Setting the Multiline Threshold Bandwidth

When using the first fit multiline select algorithm (see “Specifying the Multiline Select Algorithm” on page 8-12), you can configure a threshold of 0 to 214748364 b/s for the first available line.

By default, the threshold is set to 0; LRM uses the simple first fit algorithm. If you select a value greater than 0, LRM still uses the first available line with reservable bandwidth to service requests, but moves to the next available line after reaching the configured threshold. When all lines are at their threshold, LRM returns to using the simple first fit algorithm.

Set this parameter only if you have configured LRM on a multiline circuit and have set the Multiline Select Algorithm parameter to First Fit.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Multiline Threshold Bandwidth**[a] parameter. Click on **Help** or see the parameter description on page A-27.	
4. Click on **OK**.	You return to the Configuration Manager window.

a. You set this parameter only if you configured the Resource Manager on a multiline circuit.

Setting the Reservation Latency

By default, the maximum latency for a reserved flow packet is 50 milliseconds (ms), limiting the amount of unreserved data that the link scheduler can transmit. When the data transmit ring reaches a size where the time to transmit the data is greater than the value of this parameter, no more unreserved data will be queued.

You can set the reservation latency from 0 to 214748364 ms. Reduce the value of this parameter to obtain better delay characteristics for reserved flows, but note that overall throughput may decrease. Increase the default value to improve throughput, but note that reserved-flow delays may increase.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Reservation Latency** parameter. Click on **Help** or see the parameter description on page A-27.	
4. Click on **OK**.	You return to the Configuration Manager window.

Setting the Maximum Reservable Bandwidth for a Dataflow

By default, there is no maximum amount of bandwidth that can be reserved by a flow.

You can use the following Site Manager procedure to specify the maximum reservable bandwidth.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Largest Bandwidth** parameter. Click on **Help** or see the parameter description on page A-28.	
4. Click on **OK**.	You return to the Configuration Manager window.

Setting the Maximum Buffer Space for a Dataflow

By default, there is no maximum on the amount of buffer space that can be used by a flow.

You can use the following Site Manager procedure to specify the maximum buffer space for a flow.

Site Manager Procedure	
You do this	**System responds**
1. In the Configuration Manager window, click on any connector for a line that is configured for RSVP.	The Edit Connector window opens.
2. Click on **Edit Line Resources** in the Edit Connector window.	The Edit Line Resources window opens.
3. Set the **Largest Buffer** parameter. Click on **Help** or see the parameter description on page A-28.	
4. Click on **OK**.	You return to the Configuration Manager window.

Appendix A
Site Manager Parameters

For each parameter associated with a multicast protocol or service, this appendix provides the Site Manager menu path to each parameter, information about default settings, valid parameter options, the parameter function, instructions for setting the parameter, and the MIB object ID. Find parameters under the following topics:

Topic	Page
DVMRP Parameters	page A-2
IGMP Parameters	page A-15
IP Configuration Parameters	page A-20
Resource Manager Parameters	page A-22
OSPF Parameters for Multicast and QoS Extensions	page A-29
RSVP Parameters	page A-33

DVMRP Parameters

This section describes the DVMRP global, circuit, and tunnel parameters.

DVMRP Global Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default: Enable
Options: Enable | Disable
Function: Enables and disables DVMRP support on the router.
Instructions: To disable DVMRP once you have configured it on the router, specify Disable.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.2

Parameter: Full Update Interval

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default: 60
Options: 10 to 2000 seconds
Function: Specifies how often routing messages containing complete routing tables are sent.
Instructions: Determine the full update interval you require and specify a value.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.4

Parameter: Triggered Update Interval

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default: 5
Options: 5 or more seconds
Function: Specifies the minimum amount of time between triggered updates.
Instructions: Triggered updates are sent in the period between full updates. Issuing a full update restarts the triggered update timer. Therefore, the triggered update interval you specify must be shorter than the full update interval you specified with the Full Update Interval parameter.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.5

Parameter: Leaf Timeout

Path:	Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default:	200
Options:	25 to 4000 seconds
Function:	Specifies a value for the Leaf Timeout timer.
Instructions:	Determine the timer interval you require and specify a value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.12.1.6

Parameter: Neighbor Timeout

Path:	Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default:	140
Options:	40 to 8000 seconds
Function:	Specifies how long a connection with a router neighbor is considered active without receiving a subsequent probe or report from the neighbor.
Instructions:	Determine a neighbor timeout period and specify a value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.12.1.7

Parameter: Route Expiration Timeout

Path:	Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default:	200
Options:	20 to 4000 seconds
Function:	Specifies how long a route is considered valid without the receipt of a subsequent update indicating that the route is reachable. This value represents the duration of time that this route will be used. Upon expiration of this timer, this route is advertised as unreachable until it is refreshed or deleted.
Instructions:	Enter a value that represents the duration of time this route will be used without being refreshed.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.12.1.8

Parameter: Garbage Timeout

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Global

Default: 340

Options: 40 to 8000 seconds

Function: Specifies the duration of time that this route will be included in routing updates without the receipt of a subsequent update indicating that the route is reachable. The difference between this value and the Route Expiration Timeout value represents the duration of time that the route will be advertised as unreachable without subsequent refreshment.

Instructions: Enter a value that is greater than the value you specified for Route Expiration Timeout to allow for sufficient time for the route to be advertised as unreachable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.9

Parameter: Estimated Routes

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Global

Default: 25

Options: An integer of 10 or greater

Function: Specifies the estimated number of routes.

Instructions: Enter a value that the router can use for preallocating routing tables. For an MBone implementation, a value of 3000 or higher is recommended. Note that routes are kept on a per-source-network basis, independent of multicast groups. This number must include a route for every network that is local to a circuit configured for multicasting. This allows the router to utilize memory efficiently. Exceeding this size during router operation will not cause an error but may cause the router to consume more memory than is required.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.10

Parameter: Neighbor Probe Interval

Path:	Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default:	10
Options:	5 to 30 seconds
Function:	Specifies how often DVMRP should send a probe on interfaces from which no neighbors have been heard.
Instructions:	If your neighbor is running DVMRP *mrouted*, ensure that your probe interval value matches the value used by the neighbor.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.12.1.11

Parameter: Route Switch Timeout

Path:	Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default:	140
Options:	20 to 2000 seconds
Function:	Specifies how long DVMRP should wait, without receiving a subsequent route update from the original neighbor, before switching to a different neighbor advertising equal cost for this route.
Instructions:	If your neighbor is running DVMRP *mrouted*, the recommended value is 140 seconds.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.12.1.12

Parameter: Debug Level

Path:	Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default:	0
Options:	0 or a debug level
Function:	Turns the DVMRP log on or off for DVMRP debugging messages and specifies the level of debug messages.
Instructions:	By default (0), the log is turned off. Specify a level only when directed by Customer Support personnel to investigate a problem.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.12.1.14

Parameter: Pruning Enable

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default: Enable
Options: Enable │ Disable
Function: Enables or disables pruning functionality on the router.
Instructions: DVMRP performs route pruning by default. If you want to disable this function on the router, select Disable.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.15

Parameter: Max Routes

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default: 0
Options: 0 or an integer
Function: Specifies the number of routes that DVMRP can learn per slot. This parameter limits the number of routes that can be stored in the routing table.
Instructions: If you want DVMRP to store all learned routes in the routing table, use the default, 0. If you specify a maximum number of routes, you must include a route for every local network that has a circuit configured for multicasting.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.17

Parameter: Policy Enable

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Global
Default: Disable
Options: Enable │ Disable
Function: Enables and disables route policy functionality at the global level.
Instructions: Set this parameter as required.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.18

DVMRP Circuit Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Circuit
Default: Enable
Options: Enable | Disable
Function: Enables or disables DVMRP on this circuit.
Instructions: If you have configured DVMRP on this circuit, enter Disable to disable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.2

Parameter: Route Enable

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Circuit
Default: Enable
Options: Enable | Disable
Function: Enables or disables this circuit for routing.
Instructions: Specify Enable if you want this circuit to be used to propagate routing information, and if you want information about the source network associated with this circuit incorporated into routing updates. Specify Enable if you want multicast datagrams to be forwarded on this circuit in “native mode” -- that is, as multicast datagrams. You can configure tunnels on this circuit.

Specify Disable if you want this circuit to exist only to support unicast tunnels. If you specify Disable, all other DVMRP circuit parameters are ignored. The source network associated with this circuit is not incorporated into the routing updates.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.5

Parameter: Metric

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Circuit
Default: 1
Options: 1 to 31
Function: Specifies the cost of this interface.
Instructions: Determine the cost that you want to assign to this hop and enter a value. We recommend the following values: for a LAN, or tunnel across a single LAN, 1; for a multihop tunnel, 2 or 3; for a serial link, or tunnel across a serial link, 1; for a backup tunnel, the primary tunnel metric + 1.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.6

Parameter: Threshold

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Circuit
Default: 1
Options: 1 to 254 hops
Function: Specifies a TTL value for the interface. This value is the minimum IP TTL required for a multicast datagram to be forwarded out this interface.
Instructions: Use this parameter to control the scope of the datagrams. If the IP TTL is less than the threshold value you specify, the router drops the datagram.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.7

Parameter: Forward Cache Size

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Circuit
Default: 32
Options: 32 to 512 entries
Function: Specifies the maximum number of entries allowed in the forwarding table associated with this interface.
Instructions: When configuring a cache size, you must also keep in mind the TTL value for entries. If you specify a large cache size, a larger TTL value is recommended. For a smaller cache, a shorter TTL for cache entries is recommended.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.14

Parameter: **Forward Cache TTL**

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Circuit

Default: 700

Options: 300 to 86400 seconds

Function: Specifies the TTL for an entry in the forwarding table.

Instructions: If your neighbor is running DVMRP *mrouted,* Bay Networks recommends a TTL value of 300 seconds.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.15

Parameter: **Advertise Self**

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Circuit

Default: Advertise

Options: Advertise | Do Not Advertise

Function: Specifies whether the router advertises its own local networks over this interface.

Instructions: If you do not want routes advertised over a particular interface, set that interface to Do Not Advertise.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.16

Parameter: **Route Supply**

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Circuit

Default: Disable

Options: Enable | Disable

Function: Supplies the default route (0.0.0.0) on this interface.

Instructions: Enable this feature if you want DVMRP to advertise a default route on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.24

Parameter: Route Listen

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Circuit
Default: Disable
Options: Enable │ Disable
Function: Receives the default route (0.0.0.0) on this interface.
Instructions: Enable this feature if you want DVMRP to listen for a default route on this interface.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.25

Parameter: Report Depend Probe

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Circuit
Default: Disable
Options: Enable │ Disable
Function: Specifies the conditions under which DVMRP accepts a route report. According to the protocol, DVMRP must send a probe packet before sending the route report packet to its neighbors. Some vendor implementations, however, do not send probes before sending reports.
Instructions: By default, DVMRP processes all received route reports. Enable this feature if you want DVMRP to process only routes received from neighbors that send probe packets.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.26

DVMRP Tunnel Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Tunnel
Default: Enable
Options: Enable | Disable
Function: Enables or disables this tunnel interface.
Instructions: If you have configured this tunnel, specify Disable to disable the tunnel.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.2

Parameter: Encapsulation Mode

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Tunnel
Default: IP Datagram
Options: IP Datagram | LSSR Option
Function: Specifies whether tunneled datagrams are encapsulated within an IP datagram or loosely encapsulated using the LSSR option.
Instructions: See RFC 1075 for information about the LSSR option, which is provided for backward compatibility.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.6

Parameter: Metric

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Tunnel
Default: 1
Options: 1 to 31
Function: Specifies the cost of this tunnel.
Instructions: Determine the cost you want to assign to this tunnel and enter a value.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.7

Parameter: Threshold

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Tunnel
Default: 1
Options: 1 to 254 hops
Function: Specifies a TTL value for the tunnel. This value is the minimum IP TTL required for a multicast datagram to be forwarded out this tunnel.
Instructions: Use this parameter to control the scope of the datagrams. If the IP TTL is less than the threshold value you specify, the router drops the datagram.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.8

Parameter: Forward Cache Size

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Tunnel
Default: 32
Options: 32 to 512 entries
Function: Specifies the maximum number of entries allowed in the forwarding table associated with this tunnel interface.
Instructions: When configuring a cache size, you must also keep in mind the TTL value for entries. If you specify a large cache size, a larger TTL value is recommended. For a smaller cache, a shorter TTL for cache entries is recommended.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.16

Parameter: Forward Cache TTL

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Tunnel
Default: 300
Options: 300 to 86400 seconds
Function: Specifies the TTL for an entry in the forwarding table.
Instructions: Select an appropriate value for your configuration.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.17

Parameter: Route Supply

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP >Tunnel
Default: Disable
Options: Enable | Disable
Function: Supplies the default route (0.0.0.0) on this interface.
Instructions: Enable this feature if you want DVMRP to advertise a default route on this interface.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.25

Parameter: Route Listen

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > tunnel
Default: Disable
Options: Enable | Disable
Function: Receives the default route (0.0.0.0) on this interface.
Instructions: Enable this feature if you want DVMRP to listen for a default route on this interface.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.26

Parameter: Control Message Mode

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP >tunnel
Default: Noencaps
Options: Encaps | Noencaps
Function: Specifies the encapsulation mode for IGMP control packets.

Noencaps causes IGMP to send control messages in regular IGMP packets with the IP protocol type set to IP_PROTOCOL_IGMP.

Encaps causes IGMP to encapsulate control messages inside IP packets with the IP protocol type set to IP_PROTOCOL_IPINIP.

All DVMRP messages sent across tunnels will be encapsulated in the same way as data packets in this mode.

Instructions: Select an encapsulation mode for control messages.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.27

Parameter: Report Depend Probe

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > tunnel

Default: Disable

Options: Enable | Disable

Function: Specifies the conditions under which DVMRP accepts a route report. According to the protocol, DVMRP must send a probe packet before sending the route report packet to its neighbors. Some vendor implementations, however, do not send probes before sending reports.

Instructions: By default, DVMRP processes all received route reports. Enable this feature if you want DVMRP to process only routes received from neighbors that have sent probe packets in accordance with the protocol.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.28

Parameter: Local IP Address

Path: Configuration Manager > Protocols > IP > Multicast > DVMRP > Tunnel > Click on Add

Default: Null

Options: The unicast IP address of an interface on a circuit supporting multicasting on the local router

Function: Identifies the local end of the tunnel. To identify a unicast tunnel, you must supply the unicast IP address of both ends of the tunnel: the local interface and the remote interface.

Instructions: Use this parameter to enter the local IP address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.4

Parameter:	**Remote IP Address**
Path:	Configuration Manager > Protocols > IP > Multicast > DVMRP > Tunnel > Click on Add
Default:	Null
Options:	The unicast IP address of an interface supporting multicasting on a neighboring router
Function:	Identifies the remote end of the tunnel. To identify a unicast tunnel, you must supply the unicast IP address of both ends of the tunnel: the local interface and the remote interface.
Instructions:	Use this parameter to enter the remote IP address.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.12.3.1.5

IGMP Parameters

This section describes the IGMP global configuration, static host, and entry interface parameters.

IGMP Global Configuration Parameters

Parameter:	**Enable**
Path:	Configuration Manager > Protocols > IP > Multicast > IGMP > Global
Default:	Enable
Options:	Enable | Disable
Function:	Enables or disables this IGMP record.
Instructions:	If you configured IGMP on this router, use this parameter to disable it.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.13.1.2

Parameter: Estimated Groups

Path: Configuration Manager > Protocols > IP > Multicast > IGMP > Global

Default: 20

Options: 5 to 65535 groups

Function: Specifies the estimated number of groups that will be simultaneously active for this router.

Instructions: Determine the approximate number of groups and enter the value. This allows the router to utilize memory efficiently. Exceeding this size during router operation will not cause an error but may cause the router to consume more memory than required. The following groups are not maintained by IGMP; you do not need to include them in the count: 224.0.0.1, 224.0.0.4, 224.0.0.5, and 224.0.0.6.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.1.4

Parameter: Version Threshold Time

Path: Configuration Manager > Protocols > IP > Multicast > IGMP > Global

Default: 540

Options: 1 to 65535 seconds

Function: The number of seconds that can elapse after IGMP detects a Version 1 query before IGMP tries to become the designated querier.

Instructions: The value you specify should be greater than the Version 1 query rate of all IGMP speakers on all networks directly connected to the router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.1.5

Parameter: Debug

Path: Configuration Manager > Protocols > IP > Multicast > IGMP > Global

Default: None

Options: Receive │ Send │ Other

Function: Causes IGMP to generate the specified log messages.

Instructions: Specify zero or more options. Set this parameter to Receive to log messages received. Set the parameter to Send to log messages sent. Set the parameter to Other to log other types of IGMP messages.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.1.6

Parameter:	**Join Ack Enable**
Path:	Configuration Manager > Protocols > IP > Multicast > IGMP > Global
Default:	Disable
Options:	Enable | Disable
Function:	Indicates whether IGMP should send an immediate response (in the form of a query) to the group associated with this IGMP membership report.
Instructions:	Set this parameter as required.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.13.1.7

Parameter:	**Forward Cache Limit**
Path:	Configuration Manager > Protocols > IP > Multicast > IGMP > Global
Default:	512
Options:	64 to 65535
Function:	Specifies the maximum number of MTM forwarding cache errors.
Instructions:	Set this parameter as required.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.13.1.8

IGMP Static Host Parameters

Parameter:	**Circuit Name**
Path:	Configuration Manager > Protocols > IP > Multicast > IGMP > Global
Default:	None
Options:	A circuit number
Function:	Specifies the circuit number for the local network on which the group member resides.
Instructions:	Use the Static Group parameters to manually register the network on which the multicast group member resides.
MIB Object ID:	1.3.6.1.4.1.18.3.5.3.13.4.1.2

Parameter: **Static Group Address**

Path: Configuration Manager > Protocols > IP > Multicast > IGMP > Global

Default: None

Options: A multicast address

Function: Specifies the multicast address for the group of which the host is a member.

Instructions: Use the Static Group parameters to manually register the network on which the multicast group member resides.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.4.1.3

Parameter: **Prefix Length**

Path: Configuration Manager > Protocols > IP > Multicast > IGMP > Global

Default: None

Options: The length of the address prefix

Function: Specifies the length of the prefix of the multicast group address.

Instructions: Use the Static Group parameters to manually register the network on which the multicast group member resides.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.4.1.4

IGMP Entry Interface Parameters

Parameter: **Enable**

Path: Configuration Manager > Protocols > IP > Multicast > IGMP > Entry

Default: Enable

Options: Enable | Disable

Function: Indicates whether this IGMP interface record is enabled or disabled.

Instructions: If you configured IGMP on this interface, use this parameter to disable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.2.1.2

Parameter: **Interface Query Rate**

Path: Configuration Manager > Protocols > IP > Multicast > IGMP > Entry

Default: 120

Options: 0 to 4096 seconds

Function: Specifies how often the router sends group membership queries on the interface.

Instructions: If there are no multicast hosts on this circuit, set the parameter to 0 to disable queries. Specifying 0 affects queries only. The router still forwards multicast datagrams on this circuit. If another IGMP router on this network has taken on the query role, this router will not send queries unless it has not heard of any queries within the number of seconds specified by the Designated Router Timeout parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.2.1.5

Parameter: **Interface Membership Timeout**

Path: Configuration Manager > Protocols > IP > Multicast > IGMP > Entry

Default: 260

Options: 30 to 8192 seconds

Function: Specifies the amount of time that a local group membership is valid without the receipt of a subsequent report for that group.

Instructions: The suggested value is (2 * Query Rate) + 20.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.2.1.6

Parameter: **Designated Router Timeout**

Path: Configuration Manager > Protocols > IP > Multicast > IGMP > Entry

Default: 140

Options: 10 to 8192 seconds

Function: Specifies the amount of time that can elapse after the last host query message before the IGMP designated router is considered down.

Instructions: The value you specify should be greater than the query rate of all IGMP routers on the network.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.2.1.7

Parameter: **Max Host Response Time**

Path: Configuration Manager > Protocols > IP > Multicast > IGMP > Entry

Default: 100

Options: 1 to 100 tenths of a second

Function: Specifies a maximum value for the amount of time that a host must wait before responding to a query. IGMP places this value in the code field of an IGMP query.

Instructions: Specify a maximum response time for this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.2.1.15

Parameter: **Mtrace Entry Lifetime**

Path: Configuration Manager > Protocols > IP > Multicast > IGMP > Entry

Default: 30 seconds

Options: 30 to 8192 seconds

Function: Specifies the amount of time that a router should keep a forwarding cache entry that was created specifically for Mtrace.

Instructions: Specify an Mtrace lifetime value for the interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.2.1.18

IP Configuration Parameters

Parameter: **IP Address**

Path: Select IP from the Select Protocols window and click on OK.

Default: None

Options: 0.0.0.0 or any valid IP address

Function: Assigns a 32-bit IP address to the interface.

Instructions: Enter the IP address of the interface in dotted-decimal notation. Enter 0.0.0.0 to configure an unnumbered interface on the circuit.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.4

Parameter: Subnet Mask

Path: Select IP from the Select Protocols window and click on OK.

Default: None

Options: The Configuration Manager automatically calculates an appropriate subnet mask, depending on the class of the network to which the interface connects. However, you can change the subnet mask with this parameter.

Function: Specifies the network and subnetwork portion of the 32-bit IP address.

Instructions: Either accept the assigned subnet mask or enter another subnet mask in dotted-decimal notation. Enter 0.0.0.0 if you are configuring an unnumbered interface on the circuit.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.6

Parameter: Transmit Bcast Addr

Path: Select IP from the Select Protocols window and click on OK.

Default: 0.0.0.0

Options: 0.0.0.0 or any valid IP broadcast address

Function: Specifies the broadcast address that this IP subnet uses to broadcast packets. Accepting 0.0.0.0 for this parameter specifies that the IP router will use a broadcast address with a host portion of all 1s. Accepting 0.0.0.0 does not configure the router to use the address 0.0.0.0 to broadcast packets. For example, if you have IP address 123.1.1.1 and a subnet mask of 255.255.255.0, accepting the default value 0.0.0.0 configures the IP router to use the address 123.1.1.255 to broadcast packets. To set the explicit broadcast address of all 1s, enter 255.255.255.255 for this parameter.

Instructions: Accept the default, 0.0.0.0, unless the calculated broadcast address (host portion) of all 1s is not adequate. If this is the case, then enter the appropriate IP broadcast address in dotted-decimal notation. If you set the IP Address parameter to 0.0.0.0 (to configure an unnumbered interface), Site Manager automatically sets this parameter to 255.255.255.255.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.8

Parameter: UnNumbered Assoc Address

Path: Select IP from the Select Protocols window and click on OK.

Default: None

Options: Any valid IP address

Function: Specifies an address that IP uses when sourcing a packet. RIP uses this address to make decisions about advertising subnets over the unnumbered interface. RIP advertises subnets over the unnumbered interface if the subnets have the same mask as the associated address.

Instructions: Specify the address of any numbered interface on the router. If you are running RIP over the unnumbered interface and if you are using a subnet address as the associated address, the local and remote associated addresses should have the same network number. If you configure local and remote associated addresses using different network numbers, you must use RIP2 mode.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.110

Resource Manager Parameters

Parameter: Estimated Bandwidth

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: 0

Options: 0 to 214748364 b/s

Function: Specifies the estimated usable bandwidth for this line.

Instructions: Enter the estimated total bandwidth for this line in bits per second. To enable line resource management, enter a value greater than zero. For point-to-point lines, you can usually enter the total line speed.

For an Ethernet line, you must estimate a value, because shared-media lines may not have all of the bandwidth available. You can base the estimate on line statistics, the number of connected workstations, or other site-specific information.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.3

Parameter: Reservable Bandwidth

Path: XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: 0

Options: 0 to 214748364 b/s

Function: Specifies the reservable bandwidth for this line.

Instructions: To enable reservable resources, enter the portion of this line's bandwidth that you want to make available for RSVP requests for guaranteed service. The value must be greater than 0 and less than the value of the Estimated Bandwidth parameter. Do not make all available bandwidth reservable; leave at least 15 percent for network management and routing protocols.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.4

Parameter: Traffic Queuing Algorithm

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: None

Options: None | Priority

Function: Specifies the queuing algorithm for all reserved traffic.

Instructions: Select None to use best-effort scheduling. Use best-effort scheduling if you want to use the resource reservation protocol (RSVP) to control admission of flows to the network, but do not want to explicitly schedule each flow's packets.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.5

Parameter: Policing Algorithm

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: None

Options: None | Leaky Bucket

Function: Specifies whether or not to use the leaky bucket policing algorithm for reserved traffic. With leaky bucket policing, the Resource Manager makes sure that all RSVP packets using reserved bandwidth follow the flowspec that appears in the RSVPRSVP connect request. The Resource Manager discards any packets that do not adhere to the flowspec. Note that such policing requires additional processing by the router.

Instructions: Select None to bypass policing if an upstream router is policing traffic, or if the applications generating the reserved traffic consistently adhere to the flowspec and do not require policing.

Enter Leaky Bucket to apply a leaky bucket policing algorithm to reserved traffic.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.6

Parameter: Bandwidth Interval (Secs)

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: 10

Options: 1 to 214748364 seconds

Function: Specifies the interval over which the Resource Manager measures instantaneous bandwidth.

Instructions: Enter a time interval in seconds if the default value of 10 seconds is not sufficient.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.7

Parameter: **Inflate Reservations Percentage**

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: 0

Options: 0 to 100 percent

Function: Specifies that the Resource Manager adds a safety buffer to each RSVP reservation on this line by increasing bandwidth requests by a specified percentage.

Instructions: To inflate reservations made on this line, enter a percent value to increase each bandwidth request. Increase the value of this parameter if you notice that the Resource Manager is discarding packets because applications generating reserved traffic are exceeding their flowspecs.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.8

Parameter: **Unreserved Policing Algorithm**

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: Queue Limit

Options: Queue Limit | Leaky Bucket

Function: Specifies the policing algorithm for unreserved traffic.

Instructions: Select the Queue Limit option to restrict the number of buffers (packets) of unreserved traffic that the Resource Manager queues. This method of policing enables unreserved traffic to use available reserved bandwidth. The Unreserved Queue Length parameter specifies the maximum number of packets in the queue.

Select the Leaky Bucket option to cause the router to actively police the unreserved traffic based on the unreserved bandwidth. This method does not allow unreserved traffic to take advantage of available reserved bandwidth.

If you enable priority queuing for this line, the router uses the priority queue parameters to perform Queue Limit policing. See *Configuring Traffic Filters and Protocol Prioritization* for information about priority queue parameters.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.9

Parameter: Unreserved Queue Length

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: 20

Options: 0 to 214748364 packets (limited by available buffers)

Function: If the Unreserved Policing Algorithm parameter is set to Queue Limit, this parameter specifies the maximum number of unreserved (best-effort) packets to be held in queue for transmission.

After the queue length reaches this value, the router discards best-effort traffic when congestion occurs.

Priority queuing limits, if configured, override the value of this parameter.

Instructions: Change this value to adjust the queue length limit.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.10

Parameter: Multiline Select Algorithm

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: First Fit

Options: First Fit | Round Robin

Function: Specifies how the Resource Manager selects which line to use for a new resource request, if you have the Resource Manager configured on a multiline circuit.

Instructions: Set this parameter only if you configured the Resource Manager on a multiline circuit.

Select First Fit to always use the first available line that can service reserved traffic requests. Select Round Robin to rotate the use of lines available to service requests. All lines for a circuit must use the same algorithm. If any one line on a circuit specifies First Fit, all lines use the first fit algorithm.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.11

Parameter: Multiline Threshold Bandwidth

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: 0

Options: 0 to 214748364 b/s

Function: Specifies how the first fit algorithm works, providing you set the Multiline Select Algorithm parameter to First Fit. If you accept the default value, the Resource Manager uses the simple first fit algorithm. If you select a value greater than 0, the Resource Manager still uses the first available line with reservable bandwidth to service requests, but moves to the next available line after reaching the configured threshold. When all lines are at their thresholds, the Resource Manager resumes using the simple first fit algorithm.

Instructions: Set this parameter only if you have configured the Resource Manager on a multiline circuit and have set the Multiline Select Algorithm parameter to First Fit.

Accept the default or enter a number less than the value of the Reservable Bandwidth parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.12

Parameter: Reservation Latency

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: 50

Options: 0 to 214748364 milliseconds

Function: Specifies the maximum latency for a reserved flow packet, limiting the amount of unreserved data that the link scheduler can transmit. When the data transmit ring reaches a size where the time to transmit the data is greater than the value of this parameter, no more unreserved data will be queued.

Instructions: Reduce the value of this parameter to obtain better delay characteristics for reserved flows, but note that overall throughput may decrease. Increase the default value to improve throughput, but note that reserved-flow delays may increase.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.13

Parameter: Largest Bandwidth

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: 0 (no maximum)

Options: 0 or an integer indicating the percentage of bandwidth

Function: Specifies the maximum bandwidth that a single flow can reserve on this line.

Instructions: Set this parameter as required.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.14

Parameter: Largest Buffer

Path: Configuration Manager > XCVR, HSSI, MCT1, or COM Connector configured for the RSVP protocol > Edit Line Resources > Edit Line Resources

Default: 0 (no maximum)

Options: 0 or an integer specifying the amount of buffer space.

Function: Indicates the maximum buffer space that a single flow can reserve on this line.

Instructions: Set this parameter as required.

MIB Object ID: 1.3.6.1.4.1.18.3.5.16.2.1.3.1.15

OSPF Parameters for Multicast and QoS Extensions

OSPF Global Parameters

Parameter: Multicast Extensions

Path: Configuration Manager > Protocols > IP > OSPF > Global
Default: 0
Options: 0 (no multicast forwarding is enabled) | 1 (intra-area multicasting only) | 3 (intra-area and inter-area multicasting) | 5 (intra-area and inter-AS multicasting) | 7 (multicasting everywhere)
Function: Enables OSPF multicast extensions and specifies the type of multicasting.
Instructions: Set the parameter as required.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.19

Parameter: Multicast Deterministic

Path: Configuration Manager > Protocols > IP > OSPF > Global
Default: Nondeterministic
Options: Nondeterministic | Deterministicstrict | Deterministicloose
Function: Selects the MOSPF mode.
Instructions: Select the appropriate variation of the Dykstra algorithm.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.20

Parameter: Multicast Route Pinning

Path: Configuration Manager > Protocols > IP > OSPF > Global
Default: Nonpinned
Options: Nonpinned | Pinned
Function: Selects pinned or nonpinned mode.
Instructions: Select the appropriate variation of the Dykstra algorithm.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.21

Parameter: Opaque Capability

Path: Configuration Manager > Protocols > IP > OSPF > Global
Default: Enabled
Options: Enabled │ Disabled
Function: Controls whether or not OSPF accepts and processes opaque LSAs.
Instructions: Select Disable if you do not want OSPF to accept opaque LSAs.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.22

Parameter: Deterministic Mcast Hold Down

Path: Configuration Manager > Protocols > IP > OSPF > Global
Default: Disabled
Options: Enabled │ Disabled
Function: Controls whether or not the Hold Down feature for deterministic MOSPF is enabled.
Instructions: Enable this feature if you want a data flow to go out an interface only if there is a reservation for the flow on the interface.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.25

Parameter: Timeout Value

Path: Configuration Manager > Protocols > IP > OSPF > Global
Default: 600 seconds
Options: An integer
Function: Specifies a timer value for timing out MOSPF forward entries.
Instructions: Use the default setting.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.26

Parameter: **Max Queued Multcast Packets**

Path: Configuration Manager > Protocols > IP > OSPF > Global

Default: 64

Options: An integer indicating the number of packets

Function: Controls how many multicast packets OSPF can queue on each slot before it finishes calculating multicast trees.

Instructions: If you specify 0, OSPF drops packets before the tree is calculated.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.27

Parameter: **Dynamic TTL**

Path: Configuration Manager > Protocols > IP > OSPF > Global

Default: Disabled

Options: Enabled │ Disabled

Function: Controls whether MOSPF uses a dynamic TTL threshold. Dynamic TTL means that the TTL threshold may be different for each (source, group, downstream) tuple.

Instructions: For maximum forwarding performance, disable this feature.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.28

OSPF Interface Parameters

Parameter: Multicast Forwarding

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces

Default: Multicast

Options: Blocked | Multicast | Unicast

Function: Specifies the way multicasts should be forwarded on this interface: not forwarded, forwarded as data link multicasts, or forwarded as data link unicasts. Data link multicasting is not meaningful on point-to-point and NBMA interfaces, and setting ospfMulticastForwarding to 0 effectively disables all multicast forwarding.

Instructions: If you configured MOSPF globally, specify the way you want IP to forward multicast packets on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.30

Parameter: Opaque On

Path: Configuration Manager > Protocols > IP > OSPF > Interfaces

Default: On

Options: On | Off

Function: Controls whether or not opaque LSAs are to be flooded out this interface.

Instructions: If you have enabled the MOSPF opaque capability globally, you can turn it on and off on this interface as required.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.31

RSVP Parameters

RSVP Global Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > RSVP > Global
Default: Enable
Options: Enable | Disable
Function: Specifies the state of global RSVP.
Instructions: Disable and reenable global RSVP as required.
MIB Object ID: 1.3.6.1.4.1.18.3.5.16.1.1.1.2

Parameter: Soloist Slots

Path: Configuration Manager > Protocols > IP > RSVP > Global
Default: All slots
Options: Slots 1 to 14
Function: Specifies the slot or slots on which RSVP is eligible to run
Instructions: Select a slot or slots for global RSVP.
MIB Object ID: 1.3.6.1.4.1.18.3.5.16.1.1.1.4

Parameter: Log Filter

Path: Configuration Manager > Protocols > IP > RSVP > Global
Default: Info Log Filter
Options: Info Log Filter | Trace Log Filter | Debug Log Filter
Function: Turns the RSVP log on and off for different levels of RSVP messages.
Instructions: Select the levels for the RSVP messages you want to log.
MIB Object ID: 1.3.6.1.4.1.18.3.5.16.1.1.1.5

RSVP Interface Parameters

Parameter: Enable

Path: Configuration Manager > Protocols > IP > RSVP > Interfaces
Default: Enable
Options: Enable | Disable
Function: Sets the state of RSVP on this interface.
Instructions: Disable and reenable RSVP on the interface as required.
MIB Object ID: 1.3.6.1.4.1.18.3.5.16.1.1.6.1.2

Parameter: Refresh Blockade Multiple

Path: Configuration Manager > Protocols > IP > RSVP > Interfaces
Default: 4 intervals
Options: 1 to 65536 intervals
Function: Specifies the number of refresh intervals that RSVP will stay alive in the blockade state.
Instructions: Specify the number of refresh intervals for the blockade state.
MIB Object ID: 1.3.6.1.4.1.18.3.5.16.1.1.6.1.8

Parameter: Refresh Multiple

Path: Configuration Manager > Protocols > IP > RSVP > Interfaces
Default: 3 intervals
Options: 1 to 65536 intervals
Function: Specifies the number of refresh intervals which must elapse before a PATH or RESV message that is not being refreshed will be timed out.
Instructions: Specify the number of refresh intervals for a PATH or RESV message.
MIB Object ID: 1.3.6.1.4.1.18.3.5.16.1.1.6.1.9

Parameter: TTL

Path:	Configuration Manager > Protocols > IP > RSVP > Interfaces
Default:	0 (no TTL override)
Options:	0 to 255
Function:	Specifies a TTL value used on this interface for messages this node originates. If set to 0, the node determines the TTL via other means. This overrides the IP TTL value.
Instructions:	Set the TTL to 0 if you want the node to use the IP TTL value.
MIB Object ID:	1.3.6.1.4.1.18.3.5.16.1.1.6.1.10

Parameter: Refresh Interval

Path:	Configuration Manager > Protocols > IP > RSVP > Interfaces
Default:	3000 hundredths of a seconds (30 seconds)
Options:	An integer indicating the refresh interval in hundredths of a second
Function:	Specifies the minimum period between refresh transmissions of a given PATH or RESV message on an interface.
Instructions:	Specify the interval in hundredths of a second.
MIB Object ID:	1.3.6.1.4.1.18.3.5.16.1.1.6.1.11

Parameter: Route Delay

Path:	Configuration Manager > Protocols > IP > RSVP > Interfaces
Default:	200 hundredths of a second (2 seconds)
Options:	An integer indicating the route delay in hundredths of a second
Function:	Specifies the approximate period from the time a route is changed to the time a resulting message appears on the interface.
Instructions:	Specify the delay period in hundredths of a second.
MIB Object ID:	1.3.6.1.4.1.18.3.5.16.1.1.6.1.12

Parameter:	**Udp Required**
Path:	Configuration Manager > Protocols > IP > RSVP > Interfaces
Default:	False
Options:	True | False
Function:	Specifies whether or not UDP encapsulation is used on this interface.
Instructions:	Select True if manual configuration forces RSVP to use UDP encapsulation on the interface. Select False to disable UDP encapsulation. If you select False, RSVP uses UDP encapsulation only if it determines that a neighbor on this interface also uses UDP encapsulation.
MIB Object ID:	1.3.6.1.4.1.18.3.5.16.1.1.6.1.13

Appendix B
Multicasting Tools

This appendix describes Bay Networks Technician Interface tools that allow you to troubleshoot a multicast network. To use the tools, issue the following Technician Interface commands:

Multicast Commands	Page
mtrace	B-2
mrinfo	B-4
rsvp	B-5

mtrace

The **mtrace** command allows you trace the branch of a multicast tree from a source to a receiver for a particular multicast group. Each branch of the tree is a data path consisting of multiple hops. Mtrace displays statistics about packet rates and losses for each hop along the path.

To display Mtrace information, enter the **mtrace** command at the Technician Interface as follows:

mtrace [<*options*>] <*source*> <*receiver*> [<*group*>]

<*options*> = [**-e** <*extra_hops*>**>** | **-g** <*gateway*>**>** | **-i** <*if_addr*>**>** | **-M** | **-m** <*max_hops*>**>** | **-q** <*nqueries*>**>** | **-r** <*resp_dest*>**>** | **-s** | **-S** <*statint*>**>** | **-t** <*ttl*>**>** | **-U** | **-v** | **-w** <*wait*>**>**]

Table B-1 explains the meanings of each **mtrace** option in more detail.

Table B-1. Meanings of Mtrace Options

Option	Meaning
-e <*extra_hops*>	Mtrace tries to trace extra hops past a nonresponding router. The <*extrahops*> value is an integer.
-g <*gateway*>	Mtrace sends the trace query via a unicast packet directly to the multicast router <*gateway*> you specify.
-i <*if_addr*>	Mtrace uses <*if_address*> as the local interface address for sending the trace query and as the default address for the receiver and the response destination. The <*if_address*> is an IP address in dotted decimal notation.
-M	Mtrace requests the response using a multicast address.
-m <*max_hops*>	Mtrace sets to integer the maximum number of hops that are traced from the receiver back to the source. The default <*max_hops*> value is 32.
-q <*nqueries*>	Mtrace sets the maximum number of query attempts to integer. The default for <nqueries> value is 3.
-r <*resp_dest*>	Mtrace sends the trace response to host (that is, the response destination) rather than to the host on which Mtrace is running, or to a multicast address other than the one regisitered for this purpose (224.0.1.32).
-s	Mtrace prints a short form output including only the multicast path and not the packet rate and loss statistics.

(continued)

Table B-1. Meanings of Mtrace Options *(continued)*

Option	Meaning
-S *<statint>*	Change the interval between statistics gathering traces to interval seconds. The default *<statint>* value is 10 seconds.
-t *<ttl>*	Mtrace sets the time-to-live (TTL) for multicast trace queries and responses to the value you specify. TTL equals the number of hops. The default *<ttl>* value is 127 hops. For local queries to the "all queries" multicast gtroup, the default is 1 hop.
-U	Mtrace always requests the response using unicast rather than attempting multicast first.
-v	Mtrace uses verbose mode. In verbose mode, Mtrace shows hop times on the initial trace and statistics display. Mtrace also shows the route that was used to forward the initial trace.
-w *<wait>*	Mtrace sets the time to wait for a trace response to seconds. The default *<wait>* value is 3 seconds.

<source> is the source of the multicast tree.

<destination> is the receiver for a particular multicast group.

[*<group>*] is a multicast group. The default group is 224.2.0.1. If you do not specify a group, Mtrace does not display statistics.

For example, the following **mtrace** command specfies source 192.32.27.193, destination 192.32.199.162 and group 224.2.0.1.

```
$ mtrace 192.32.27.193 192.32.199.162 224.2.0.1
Mtrace from 192.32.27.193 to 192.32.199.162 via group 224.2.0.1
Querying full reverse path...
  0  192.32.199.162
 -1  192.32.199.162  DVMRP  thresh^ 1
 -2  192.32.27.193
Round trip time 1039 ms; total ttl of 2 required.

Waiting to accumulate statistics...
  Source       Response Dest    Overall    Packet Statistics For Traffic
From
192.32.27.193   192.32.199.161   Packet      192.32.27.193 To 224.2.0.1
     v       __/  rtt    3ms     Rate        Lost/Sent = Pct  Rate
192.32.27.193
192.32.199.162
     v      \__   ttl    2         0 pps
192.32.199.162     * * *
  Receiver      Query Source
```

mrinfo

The **mrinfo** command displays the capabilities of a DVMRP multicast router. It shows whether the router supports mtrace and pruning, provides revision information, and provides the link characteristics for every link on the router.

To display Mrinfo information, enter the **mrinfo** command at the Technician Interface as follows:

mrinfo [<*options*>] <*router_address*>

<*options*> = [< **-r** <*retry_count*> | **-t** <*timeout_count*>]

Table B-2 explains the meanings of each **mrinfo** option in more detail.

Table B-2. Meanings of Mrinfo Options

Option	Meaning
-r <*retry_count*>	Mrinfo sets the neighbor query retry limit to <*retry_count*>. The default is 3 retries.
-t <*timeout_count*>	Mrinfo sets the number of seconds to wait for neighbor query reply to <*timeout_count*>. The default timeout is 3 seconds.

<*router_address*> is the address of the DVMRP router whose capabilities you want to display.

For example, the following **mrinfo** command displays the capabilities of DVMRP router 192.32.27.46. The DVMRP router is running version 3.255 and supports pruning and Mtrace.

```
$ mrinfo 192.32.27.46
192.32.27.46 [version 3.255,prune,genid,mtrace]:
  192.32.27.46 -> 192.32.27.39 [1/1]
  192.32.27.46 -> 192.32.27.47 [1/1]
  192.32.27.46 -> 192.32.191.2 [1/1/tunnel]
  192.32.27.161 -> 192.32.27.162 [1/1]
```

rsvp

The **rsvp** command allows you to view information about RSVP and set global configuration parameters.

Enter the **rsvp** command at the Technician Interface as follows:

rsvp *<subcommand>* [*<options>*]

<subcommand> = <**session** | **phop** | **oi** | **path** | **resv**>

Table B-3 explains the meanings of each **rsvp** subcommand in more detail.

Table B-3. IP Subcommand Meanings

Subcommand	System Displays
session	Known RSVP sessions
phop	Known RSVP previous hops
oi	Known RSVP outgoing interfaces
path	The RSVP path state
resv	The RSVP RESV state
blockade	The RSVP blockade state

<options> = [**-d** *<destination IP>* [*<protocol ID>*|*<destination port>*] | **-p** *<previous hop>* | **-f** *<source IP>* [*<source port>*] | **-n** *<next hop>*]

The *<options>* apply to subcommands, as described in Table B-4.

Table B-4. Option Descriptions

Option	Effect	Applicable Subcommands
-d *<destination IP>* [*<protocol ID>* | *<destination port>*]	Retrieves RSVP data for the session you specify	session | phop | oi | path | resv
-p *<previous hop>*	Retrieves only sessions with a previous hop of <previous hop>	phop | path | resv
-f *<source IP>* [*<source port>*]	Retrieves only oi, path, or resv information for sessions with a source address of <source IP> and, optionally, a source port of <source port>	oi | path | resv
-n *<next hop>*	Retrieves only reservation request information with a next hop of <next hop>	resv

For example, the following **rsvp** command specifies the RSVP previous hops:

```
$ rsvp phop

RSVP Previous Hops:

       Session
IP dest/protocol/dest port   Previous Hop
--------------------------   --------------
225.1.1.1 17 5000            199.1.4.5
225.1.1.1 17 5000            200.1.1.212
225.1.1.1 17 5000            201.1.1.212

   3 previous hops matched
```

Index

A

addresses
  multicast, 1-3
  supplying for tunnel, 4-34
Advertise Self parameter, 4-29, A-9

B

Bandwidth, A-24
Bandwidth Interval (Secs) parameter, 8-8, A-24
Bay Networks Press, xviii
border router (BR) in multicasting, 5-2
boundary router (ASBR) in multicasting, 5-2

C

Circuit Name parameter, 3-20, A-17
Control, A-13
Control Message Mode parameter, 4-43, A-13
customer support
  programs, xix
  Technical Solutions Centers, xix

D

Debug Level parameter, 4-19, A-5
debug messages, logging, 4-19
Debug parameter, 3-9, A-16
default route, DVMRP circuit, listening for, 4-31
Delay, A-35
Designated Router Timeout parameter, 3-17, A-19
Deterministic Mcast Hold Down parameter, 6-3, A-30
deterministic mode for MOSPF, 5-5

disabling
advertisement of local networks, 4-29
deterministic MOSPF, 5-5
DVMRP circuit, 4-22
DVMRP global, 4-9
IGMP global, 3-5
IGMP interface, 3-13
MOSPF, 5-4
multicast support on unicast tunnel, 4-23
RSVP on an interface, 7-7
tunnels, 4-35
UDP encapsulation, 7-13
DVMRP
accepting report from nonstandard neighbor, 4-44
calculating route metric, 4-24
choosing a route, 4-5
comparing route costs, 4-24
cost of a tunnel, 4-37
enabling and disabling, 4-9
advertisement of local networks, 4-29
circuit, 4-22
encapsulation mode for multicast datagrams, 4-36
estimating number of routes for, 4-18
forwarding table size, 4-27
listening for default route, 4-31
logging debug messages, 4-19
number of routes, 4-20
probe interval, 4-16
route switch interval, 4-17
threshold value of tunnel, 4-38
tree pruning, 4-11
Dynamic TTL parameter, 5-8, A-31

E

Enable parameter
DVMRP
circuit, 4-22, A-7
global, 4-9, A-2
tunnel, 4-35, A-11
IGMP
global, 3-5, A-15
interface, 3-13, A-18
RSVP, 7-7
enabling
advertisement of local networks, 4-29
deterministic MOSPF, 5-5
DVMRP circuit, 4-22
DVMRP global, 4-9
IGMP global, 3-5
IGMP interface, 3-13
MOSPF, 5-4
MOSPF on an OSPF interface, 5-9
multicast support on unicast tunnel, 4-23
RSVP on an interface, 7-7
tunnels, 4-35
UDP encapsulation, 7-13
Encapsulation Mode parameter, 4-36, A-11
Estimated Bandwidth parameter, 2-5, 8-4
Estimated Groups parameter, 3-7, A-16
Estimated Routes parameter, 4-18, A-4

F

First Fit multiline select algorithm, 8-12
Forward Cache Limit parameter, 3-11, A-17
Forward Cache Size parameter, 4-27, 4-39, A-8, A-12
Forward Cache TTL parameter, 4-28, 4-40, A-9, A-12
forwarding table, defined, 4-27
Full Update Interval parameter, 4-10, A-2

G

Garbage Timeout parameter, 4-13, A-4

H

host groups, types of, 1-2

I

IGMP (Internet Group Management Protocol), 1-3
  adding a static host entry, 3-20
  designated router timeout interval, 3-17
  enabling and disabling
    global, 3-5
    interface, 3-13
  estimating number of groups, 3-6
  host leave message, 3-3
  host repots, 3-3
  lifetime of a cache entry for Mtrace, 3-19
  logging, 3-9
  maximum host response time, 3-18
  membership timeout interval, 3-16
  queries, 3-2
  query rate, 3-14
  version threshold time, 3-8

IGMP (internet Group Management Protocol)
  function of, 3-2

Inflate, A-25

Inflate Reservations Percentage parameter, 8-9, A-25

Interface Membership Timeout parameter, 3-16, A-19

Interface Query Rate parameter, 3-15, A-19

internal router (IR) in multicasting, 5-2

Internet Group Management Protocol
  See IGMP (Internet Group Management Protocol)

IP Address parameter, A-20
  IP Configuration
    IGMP and DVMRP, 2-3
    RSVP, 2-5

J

Join Ack Enable parameter, 3-10, A-17

L

Largest Bandwidth parameter, 8-15, A-28

Largest Buffer parameter, 8-16, A-28

Leaf Timeout parameter, 4-15, A-3

leaky bucket policing algorithm, 8-10

Line Resource Manager
  policing algorithm, 8-7, 8-10

Line Resource Manager (LRM)
  bandwidth interval, 8-8
  inflating reservations, 8-9
  latency for reserved flow packet, 8-14
  maximum number of unreserved packets in queue, 8-11
  multiline select algorithms, 8-12
  multiline threshold bandwidth, 8-13
  overview, 8-2
  traffic queuing algorithm, 8-6

Local IP Address parameter, 4-34, A-14

Log, A-33

Log Filter parameter, 7-6, A-33

logging
  IGMP messages, 3-9
  RSVP messages, 7-6

logging, debug messages, 4-19

LSSR option, 4-36

M

Max Host Response Time parameter, 3-18, A-20

Max Queued Multcast Packets parameter, 5-7, A-31

Max Routes parameter, 4-20, 4-21

Metric parameter, 4-24, 4-37, A-8, A-11

MOSPF (Multicasting Extensions to OSPF), 1-4
deterministic hold-down feature, 6-3
deterministic mode, 5-5
enabling and disabling, 5-4
enabling on an OSPF interface, 5-9
forwarding timeout value, 6-6
nondeterministic mode, 5-5
opaque capability
enabling on an interface, 6-6
globally enabling, 6-5
overview, 5-1
route pinning, 6-4
starting, 2-4
types of multicasting support, 5-3
Mtrace Entry Lifetime parameter, 3-19, A-20
Mtrace utility, 3-19
Multicast Deterministic parameter, 5-5, A-29
Multicast Extensions parameter, 2-4, 5-4, A-29
Multicast Forwarding parameter, 5-9, A-32
Multicast Route Pinning parameter, 6-4, A-29
multicasting
addresses for, 1-3
aging a route, 4-12
border router (BR), 5-2
boundary router (ASBR), 5-2
comparing routes, 4-7
creating a shortest path tree, 4-7
IGMP, 1-3
internal router (IR), 5-2
leaf network, 4-7
threshold, 4-25
tunnel, 4-2
types supported by MOSPF, 5-3
Multicasting Extensions to OSPF
See MOSPF
Multiline Select Algorithm parameter, 8-12, A-26
Multiline Threshold Bandwidth parameter, 8-13, A-27

N

Neighbor Probe Interval parameter, 4-16, A-5
Neighbor Timeout parameter, 4-14, A-3
neighbors
nonstandard, 4-44
nondeterministic mode for MOSPF, 5-5

O

Opaque Capability parameter, 6-5, A-30
Opaque On parameter, 6-6, A-32

P

permanent host group, 1-2
Policing Algorithm parameter, 8-7, A-24
Prefix Length parameter, 3-20, A-18
priority queuing, 8-11, A-26
Pruning Enable parameter, 4-11, A-6
pruning the shortest-path tree, 4-11
publications, ordering, xviii

Q

QOSPF (Quality of Service Path First)
adding extensions to OSPF, 6-2
functions, 6-1
Quality of Service (QoS), 8-2
queue limit policing algorithm, 8-10

R

Refresh, A-35
Refresh Blockade Multiple parameter, 7-10, A-34
Refresh Interval parameter, 7-8, A-35
Refresh Multiple parameter, 7-9, A-34
Remote IP Address parameter, 4-34, A-15
Report Depend Probe parameter, 4-44, A-10
Reservable Bandwidth parameter, 2-5, 8-5
Reservation Latency parameter, 8-14, A-27
Resource Reservation Protocol
See RSVP (Resource Reservation Protocol)

RFC 1075, 4-36
Round Robin multiline select algorithm, 8-12
Route Delay parameter, 7-12, A-35
Route Depend Probe parameter, 4-32, A-14
Route Enable parameter, 4-23, A-7
Route Expiration Timeout parameter, 4-12, A-3
Route Listen parameter, 4-31, 4-42, A-10
Route Supply parameter, 4-30, 4-41, A-9
Route Switch Timeout parameter, 4-17, A-5
RSVP (Resource ReSerVation Protocol)
  estimating bandwidth, 8-4
  setting reservable bandwidth, 8-5
  starting, 2-5
RSVP (Resource Reservation Protocol)
  configuring with the Technician Interface, 1-4
  enabling and disabling, 7-7
  function, 7-1
  lifetime multiplier, 7-9
  loading on a slot, 7-5
  logging messages, 7-6
  refresh blockade multiplier, 7-10
  refresh timer, 7-8
  route delay value, 7-12
  TTL override, 7-11
  UDP encapsulation, 7-13

S

Soloist Slots parameter, 7-5
source network
  defined, 4-4
starting
  DVMRP, 2-3
  MOSPF, 2-4
  RSVP, 2-5
Static Group Address parameter, 3-20, A-18
Subnet Mask parameter, IP Configuration
  IGMP and DVMRP, 2-3
  RSVP, 2-5

T

Technical Solutions Centers, xix
Threshold parameter, 4-26, 4-38, A-8, A-12
threshold values, function of, 4-25
Timeout Value parameter, 5-6, A-30
timer
  garbage, 4-13
  leaf, 4-15
  neighbor report, 4-14
  route expiration, 4-12
  RSVP default refresh, 7-8
traffic filters, affect on multicast data packets, 1-2
Traffic Queuing Algorithm parameter, 8-6, A-23
transient host group, 1-2
Transmit Bcast Addr parameter, IP Configuration
  IGMP and DVMRP, 2-3
  RSVP, 2-5
tree pruning, 4-11
Triggered Update Interval parameter, 4-10, A-2
triggered updates, specifying the time between, 4-10
TTL parameter, 7-11

tunnels
addresses for, 4-34
connecting neighbors, 4-2
cost, 4-37
enabling and disabling, 4-35
forwarding table for, 4-39
threshold value for, 4-38
unicast, disabling multicast support on, 4-23

U

Udp Required parameter, 7-13
UnNumbered Assoc Address parameter, IP Configuration
RSVP, 2-5
UnNumbered Assoc Address, IP Configuration
IGMP and DVMRP, 2-3
Unreserved Queue Length parameter, 8-11, A-26
UPD encapsulation, 7-13
updates
setting the interval, 4-10
triggered, 4-10

V

Version Threshold Time parameter, 3-8, A-16